Contents

Part A

Implicational Calculus

Chapter 1

Syntax

In Part A we study the implicational fragment of Propositional Calculus, in various formulations. Different approaches not only provide a better understanding of the concept of implication: their equivalence proofs are instructive, and will introduce useful technical tools and notions in a simplified environment, with technicalities reduced to a minimum. In following chapters the treatment will be expanded to cover increasingly more powerful systems.

We consider *Natural Deduction* as the basic system, and reading of this chapter can be confined to Section 1. We see the remaining approaches as providing alternative formulations of the basic system, with different aims: to *synthesize* its proofs (Section 2), and to set up a calculus to *describe* them (Section 3). Parts B and C will introduce still other approaches, in terms of *categories* and *λ-calculus*.

Implicational Calculus

The **language** consists of:

- propositional letters p, q, r, ...
- parentheses '(' and ')'
- the connective $\rightarrow$ (*implication*).

Formulas are defined inductively as follows:

- propositional letters are formulas
- if α and β are formulas, so is $(\alpha \rightarrow \beta)$.

To increase readability some parentheses can either be omitted, when no confusion arises, or written differently, e.g. as '[' and ']'. We will use lowercase Roman letters such as p for propositional letters, and lowercase Greek letters such as α for formulas.

The main goal of this chapter is to determine which of the formulas of the Implicational Calculus can be considered 'true'. This of course requires a presupposed meaning of the connective →, which is intuitively taken as representing 'implication'. We will introduce various different but equivalent analyses to make this intuitive meaning explicit and precise.

1.1 Natural Deduction

The system of Natural Deduction is based on rules that show how to continue a given proof from assumptions. Proofs will take the form of a tree, whose leaves are (at the moment of their creation) *assumptions* (some of which can later be discharged), and whose roots are *conclusions*.

When we want to indicate explicitly that a proof tree $\mathcal{D}$ has undischarged assumptions in the set Γ and a conclusion β, we will use the following representation:

$$\begin{array}{c}\Gamma\\ \mathcal{D}\\ \beta.\end{array}$$

Proof trees

In our first approach to implication we try to capture the usual mathematical practice: to *prove* $\alpha \to \beta$, we prove β under the assumption α. A subtle distinction occurs here: while we are trying to reach β from α, the latter is seen and used as an assumption; but after we reached β, a change in the status of α has occurred, since we have proved the wanted implication $\alpha \to \beta$, which already contains the information that α was used as an assumption. Technically, we say that α has been used as a temporary assumption and then *discharged* at the end. We can picture the situation graphically in full generality (using a finite set Γ of additional assumptions), as the step

$$\text{from} \quad \begin{array}{c}\Gamma, \alpha\\ \mathcal{D}_\beta\\ \beta\end{array} \quad \text{to} \quad \dfrac{\begin{array}{c}\Gamma, [\alpha]^{(1)}\\ \mathcal{D}_\beta\\ \beta\end{array}}{\alpha^{(1)} \to \beta.}$$

Here $\mathcal{D}_\beta$ is a proof of β from the set of assumptions $\Gamma \cup \{\alpha\}$ (written as Γ, α to increase readability). The square brackets around α indicate the discharge, and are needed to keep track of which assumptions are still active in a proof. The number (1) is optional, and it has two functions: in the bottom part of the proof it tells *when* the assumption is discharged, and in the top part it shows *which* occurrence(s) of α have been discharged. Without such a numbering, it would be impossible to understand how a proof has been constructed. For example, to distinguish between

$$\dfrac{\dfrac{[\alpha]^{(1)}}{\alpha^{(1)} \to \alpha}}{\alpha^{(2)} \to (\alpha \to \alpha).} \qquad \text{and} \qquad \dfrac{\dfrac{[\alpha]^{(2)}}{\alpha^{(1)} \to \alpha}}{\alpha^{(2)} \to (\alpha \to \alpha).}$$

Note that there is no occurrence of α in the first proof that corresponds to the second discharge (numbered with 2). This is in accord with the intuition that if we have a proof of β that does not use a certain assumption, we can think of that assumption as sitting idle at the beginning of the proof, and discharge it without having actually used it. In particular, *we may discharge an assumption that has no occurrence in a proof*. An equivalent but more explicit way of picturing the same proof would be:

$$\dfrac{\dfrac{[\alpha]^{(1)} \qquad [\alpha]^{(2)}}{\alpha^{(1)} \to \alpha}}{\alpha^{(2)} \to (\alpha \to \alpha).}$$

The second proof above illustrates a more complicated situation: we not only discharge in the first step an ‘occurrence’ of α which is not there, but do not discharge an occurrence of α which is present, keeping it for later use. This is in accord with the intuition that different occurrences of an assumption in the same proof may serve different purposes, and we should have the possibility of distinguishing among them. In particular, *we may discharge only some of the occurrences of an assumption in a proof*.

To sum up, *we may discharge none, some, or all occurrences of an assumption in a proof*. This shows that it is not really single *formulas* that are treated as assumptions, but rather *packets of occurrences* of formulas. We can simply say that what is really discharged is every occurrence of a formula in a single packet of assumptions. Occurrences of the same formula in different packets are treated as different formulas, that happen to coincide. Thus, for example,

$$\dfrac{\dfrac{[\alpha]^{(1)} \qquad [\alpha]^{(2)}}{\alpha^{(1)} \to \alpha}}{\alpha^{(2)} \to (\alpha \to \alpha)}$$

is actually a special case of the following, with $\alpha = \beta$:

$$\dfrac{\dfrac{[\alpha]^{(1)} \qquad [\beta]^{(2)}}{\alpha^{(1)} \to \beta}}{\beta^{(2)} \to (\alpha \to \beta).}$$

In particular, when saying that Γ is the set of *assumptions* of a given proof, we really mean a set of *packets of occurrences of assumptions*.

We not only want to be able to prove implications, but also to *use* them. The way to do this is suggested by a venerable principle already known to the Greeks, the socalled *Modus Ponens*: this is the principle that, since a proof $\mathcal{D}_{\alpha\to\beta}$ of $\alpha \to \beta$ codes a proof $\mathcal{D}_\beta$ of β with α used as an assumption, it can be completed to a proof of β by the addition of a proof $\mathcal{D}_\alpha$ of α. We can picture the situation graphically in full generality (using a finite set Γ of additional assumptions), as the step

$$\text{from}\quad \begin{array}{c}\Gamma\\ \mathcal{D}_{\alpha\to\beta}\\ \alpha\to\beta\end{array} \quad\text{and}\quad \begin{array}{c}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{array}\quad\text{to}\quad \dfrac{\begin{array}{cc}\Gamma & \Gamma\\ \mathcal{D}_{\alpha\to\beta} & \mathcal{D}_\alpha\\ \alpha\to\beta & \alpha\end{array}}{\beta.}$$

We now have the rules of Natural Deduction for implication: proof trees can be constructed inductively, by starting from a set of assumptions divided into (possibly empty) packets of occurences, and by continuing them by the rules of introduction or elimination of →.

The following is a non trivial example of proof:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{[\alpha]^{(1)} \quad [\alpha\to\beta]^{(3)}}{\beta} \qquad [\beta\to\gamma]^{(2)}}{\gamma}}{\alpha^{(1)}\to\gamma}}{(\beta\to\gamma)^{(2)}\to(\alpha\to\gamma)}}{(\alpha\to\beta)^{(3)}\to[(\beta\to\gamma)\to(\alpha\to\gamma)].}$$

Consequence relation

We now give a formal inductive definition of the consequence relation $\vdash_{\mathcal{N}}$. The intuitive meaning of $\Gamma \vdash_{\mathcal{N}} \beta$ is that Γ is a finite set of packets of assumptions from which β can be deduced by a proof of the kind described above.

Definition 1.1.1 (Gentzen [1935]) *The relation $\vdash_{\mathcal{N}}$ is inductively defined as follows.*

1. **Assumptions**. *An assumption can be deduced from a set of assumptions to which it belongs:*

$$\Gamma, \beta \vdash_{\mathcal{N}} \beta.$$

2. **→-Introduction**. *If β is deducible from Γ and a packet of assumptions of α, then $\alpha \to \beta$ is deducible from Γ, by discharging that packet:*

$$\frac{\Gamma, \alpha \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \to \beta.}$$

3. **→-Elimination**. *If α and $\alpha \to \beta$ are both deducible from Γ, then so is β:*

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma \vdash_{\mathcal{N}} \alpha \to \beta}{\Gamma \vdash_{\mathcal{N}} \beta.}$$

In an application of the →-elimination rule, $\alpha \to \beta$ is called the *major premise* of the rule, and α the *minor premise*.

The formulation of →-introduction in the form

$$\frac{\Gamma, \alpha \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \to \beta.}$$

can be read as: 'α informally implies β' implies 'α formally implies β'. This separates three different meanings of the intuitive concepts of implication: one internal to $\mathcal{N}$, and two different ones external to $\mathcal{N}$. The first is captured by the connective →, and is expressed by $\alpha \to \beta$. The second is captured by the metalinguistic symbol $\vdash_{\mathcal{N}}$, that defines $\mathcal{N}$ but does not belong to it, and is expressed by $\alpha \vdash_{\mathcal{N}} \beta$, i.e. by the fact that β can be deduced in $\mathcal{N}$ from α. The third has nothing to do with $\mathcal{N}$, and is the informal meaning of implication used (metalinguistically) in mathematical reasonings about the linguistic objects of discourse, in our case derivations and proofs in $\mathcal{N}$: in the rule above, it is indicated by the horizontal line separating the premiss from the conclusion.

Actually, the →-introduction and →-elimination rules together establish that the *two concepts of implications in $\mathcal{N}$ coincide*, in the sense that

$$\Gamma, \alpha \vdash_{\mathcal{N}} \beta \text{ if and only if } \Gamma \vdash_{\mathcal{N}} \alpha \to \beta.$$

The left to right direction is simply a restatement of →-introduction. The right to left direction follows from →-elimination, because if we have $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ then there is a derivation $\mathcal{D}$ in $\mathcal{N}$ of $\alpha \to \beta$ from Γ, and thus the following is a derivation of β from Γ and α:

$$\frac{\begin{array}{c}\Gamma \\ \mathcal{D} \\ \alpha \to \beta\end{array} \qquad \alpha}{\beta.}$$

Hence $\Gamma, \alpha \vdash_{\mathcal{N}} \beta$.

The following is a translation of the example of proof given at the end of the previous subsection, where $\Gamma = \{\alpha, \alpha \to \beta, \beta \to \gamma\}$:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma \vdash_{\mathcal{N}} \alpha \to \beta}{\Gamma \vdash_{\mathcal{N}} \beta} \qquad \Gamma \vdash_{\mathcal{N}} \beta \to \gamma}{\alpha, \alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{N}} \gamma}}{\alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{N}} \alpha \to \gamma}}{\alpha \to \beta \vdash_{\mathcal{N}} (\beta \to \gamma) \to (\alpha \to \gamma)}}{\vdash_{\mathcal{N}} (\alpha \to \beta) \to [(\beta \to \gamma) \to (\alpha \to \gamma)].}$$

This is actually the same proof as before, except for the use of different devices to keep track of premises. The translation is obtained by forgetting about the labels of formulas, and by writing on the left of $\vdash_{\mathcal{N}}$ all premises not yet discharged. Conversely, if we erased the left-hand-side of $\vdash_{\mathcal{N}}$ in the translation, and labeled the premises that are moved from left to right by →-introduction, we would get back the original proof.

Since the two formalisms are equivalent, we will use the most convenient one in each practical case.

Normal proofs

The →-introduction rule defines a proof of $\alpha \to \beta$ as an incomplete proof $\mathcal{D}_\beta$ of β from a packet of assumptions α, waiting for a completion. The →-elimination rule allows the completion of such a proof, whenever we have a proof $\mathcal{D}_\alpha$ of α. Taken together, the two rules combine in transforming the two proofs $\mathcal{D}_\alpha$ and $\mathcal{D}_\beta$ into a proof of β, as follows:

$$\dfrac{\dfrac{\begin{array}{c}\Gamma, [\alpha]\\ \mathcal{D}_\beta\\ \beta\end{array}}{\alpha \to \beta} \qquad \begin{array}{c}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{array}}{\beta.} \tag{1.1}$$

A more direct way of getting from $\mathcal{D}_\alpha$ and $\mathcal{D}_\beta$ to a proof of β would be the following:

$$\begin{array}{c} \quad\Gamma\\ \quad\mathcal{D}_\alpha\\ \Gamma, \quad \alpha\\ \mathcal{D}_\beta\\ \beta, \end{array} \tag{1.2}$$

i.e. to directly substitute $\mathcal{D}_\alpha$ above every occurrence of α in the packet of assumption used in $\mathcal{D}_\beta$.

The first approach is of course useful in mathematical practice, since we usually get the proofs $\mathcal{D}_\alpha$ and $\mathcal{D}_\beta$ at different times, and the proof of $\mathcal{D}_{\alpha\to\beta}$ obtained from $\mathcal{D}_\beta$ by →-introduction records, for future use, the fact that we have obtained an incomplete proof of β from the assumption α. After we have also obtained the proof $\mathcal{D}_\alpha$, a single additional step merges it with $\mathcal{D}_{\alpha\to\beta}$ to get a proof of β.

However, the second approach has the advantage of being more direct, and of not going through the unnecessary detour of a →-introduction followed by a →-elimination. The occurrence of $\alpha \to \beta$ in 1.1 is called a **maximum** relative to →, and the step from 1.1 to 1.2 is called a **maximum elimination**. A proof is in **normal form** if it has no maxima.

The proof considered above is an example of a normal proof:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{[\alpha]^{(1)} \quad [\alpha \to \beta]^{(3)}}{\beta} \qquad [\beta \to \gamma]^{(2)}}{\gamma}}{\alpha^{(1)} \to \gamma}}{(\beta \to \gamma)^{(2)} \to (\alpha \to \gamma)}}{(\alpha \to \beta)^{(3)} \to [(\beta \to \gamma) \to (\alpha \to \gamma)].}$$

As we can see, the proof is nicely divided into an upper part consisting of eliminations, and a lower part consisting of introductions.

This division is made possible by the fact that in a normal proof there is no introduction followed by an elimination of the same arrow, and it should hold in general. The only difficulty is how to express the property appropriately, since an introduction *could* be followed by an elimination, although not on the same arrow. In other words, there could be an *insignificant* alternation of →-introduction and →-elimination, in the sense that one → is first introduced in the minor premise, and then a *different* → is eliminated in the major premise, as in:

$$\dfrac{\dfrac{\begin{array}{c}\alpha \\ \mathcal{D} \\ \beta\end{array}}{\alpha \to \beta} \qquad (\alpha \to \beta) \to \gamma}{\gamma.}$$

To take care of this possibility, we introduce the notion of a *descending path*, on which there is indeed a separation between an upper and a lower part.

More precisely, in a normal proof of α from assumptions Γ a **descending path** is a branch of the proof tree starting from a leaf (i.e. either a formula in Γ or a discharged assumption) and proceeding through either →-introductions or major premises of →-eliminations, until either the conclusion α or the minor premise of a →-elimination rule is reached. Briefly, a descending path is a branch of the proof tree not going through a minor premise of a →-elimination.

Proposition 1.1.2 Structure of Normal Proofs (Prawitz [1965]) *For a normal proof of $\mathcal{N}$ the following hold:*

1. **Elimination-Introduction Separation**. *Any descending path consists of two (possibly empty) parts: a first (upper) one going only through →-eliminations, and a second (lower) one going only through →-introductions.*

2. **Subformula Property**. *Any formula occurring in the proof is a subformula of either an undischarged assumption or the conclusion.*

Proof. Since a descending path goes through either →-introductions or major premises of →-eliminations, the only way a →-introduction on it can be followed by a →-elimination is when the major premise of the latter is a maximum. But no maximum exists, since the proof is normal. This proves part 1.

To prove part 2, notice that if there was a counterexample to the subformula property, there would be one of maximal length (in terms of symbols). For such a counterexample, there are two possible cases:

- If it occurs in a →-elimination rule

  $$\frac{\alpha \qquad \alpha \to \beta}{\beta,}$$

  it can be neither α nor β, otherwise $\alpha \to \beta$ would be a counterexample of greater length, and hence it must be $\alpha \to \beta$. But it cannot have been introduced before, otherwise it would be a maximum and the proof would not be normal. And it cannot be discharged afterwards, otherwise it would be the premise of a counterexample of greater length. Then it must be an undischarged assumption, and it is not a counterexample.

- If it occurs in a →-introduction rule

  $$\frac{\begin{array}{c}\alpha\\ \mathcal{D}\\ \beta\end{array}}{\alpha \to \beta,}$$

  it must be $\alpha \to \beta$ as in the first case. But it cannot be eliminated afterwards, otherwise it would be a maximum and the proof would not be normal. And it cannot be part of a following introduction, otherwise it would be the conclusion of a counterexample of greater length. Then it must be the conclusion, and it is not a counterexample. □

A Normalization procedure

Having shown that proofs in normal form have nice properties, we now prove that every proof can be reduced to one in normal form.

Theorem 1.1.3 Weak Normalization (Prawitz [1965]) *Every proof in the Natural Deduction system can be transformed into a normal proof, by means of an appropriate sequence of maxima eliminations.*

Proof. Notice that the elimination of a maximum

$$\dfrac{\dfrac{\begin{array}{c}\Gamma,\alpha\\ \mathcal{D}_\beta\\ \beta\end{array}}{\alpha\to\beta} \qquad \begin{array}{c}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{array}}{\beta} \qquad \text{into} \qquad \begin{array}{c}\Gamma\\ \mathcal{D}_\alpha\\ \Gamma,\ \ \alpha\\ \mathcal{D}_\beta\\ \beta,\end{array}$$

in a proof $\mathcal{D}$ can have the following two bad effects:

- it can increase the total number of maxima, since it reproduces $\mathcal{D}_\alpha$ (and hence all maxima occurring in it) above every occurrence of α in the package of assumptions used in $\mathcal{D}_\beta$, and there may be many such occurrences;

- it can introduce new maxima, in two different ways:

  - if $\alpha = \gamma \to \delta$ and $\mathcal{D}_\alpha$ ends with a →-introduction, by turning into a maximum every occurrence of α below which $\mathcal{D}_\beta$ continues with a →-elimination
  - if $\beta = \gamma \to \delta$, $\mathcal{D}_\beta$ ends with a →-introduction, and $\mathcal{D}$ continues below β with a →-elimination, by turning into a maximum that occurrence of β.

The main observation is that the second obstacle is not traumatic, since the new maxima α or β possibly introduced are of complexity lower than the one $\alpha \to \beta$ being eliminated. The appropriate measure of complexity is in this case the **degree** of a formula, defined inductively as follows:

- propositional letters have degree 0

- the degree of $\alpha \to \beta$ is 1 plus the the greatest of the degrees of α and β.

The idea of the normalizing procedure is thus to eliminate, at every step, a maximum of greatest degree, until all of them have been disposed of. The first obstacle is overcome by choosing, at every step, *a maximum $\alpha \to \beta$ of greatest degree, such that in $\mathcal{D}_\alpha$ no maximum of greatest degree occurs* (so that only the number of maxima of degree smaller than the greatest one can be increased).

By so doing, at every step we eliminate one maximum of greatest degree, and do not introduce new ones of the same degree. Once the last maximum of greatest degree has been eliminated, we attack the ones of the next greatest degree (whose number, in the meantime, may have greatly increased), and so on, until all maxima have been eliminated. □

Formally, the proof of the Weak Normalization Theorem is by socalled ω^2-*induction*, i.e. induction on pairs of natural numbers (a, b) lexicographically ordered by

$$(a, b) \prec (a', b') \iff (a < a') \vee (a = a' \wedge b < b').$$

Indeed, at every step the pair

$$(\text{greatest degree}, \text{number of maxima of greatest degree})$$

strictly decreases in the ordering $\prec$ (i.e., either the greatest degree decreases, or it remains the same but the number of maxima with greatest degree decreases by one).

1.2 Hilbert Systems

The common mathematical practice makes use of socalled **Hilbert systems**, in which the notion of *theorem* is inductively defined as follows:

- axioms are theorems
- formulas deduced from theorems by the use of a deduction rule (in Propositional Calculus: only Modus Ponens) are theorems.

Consequence relation

We can extend the notion of theorem to that of *formula deduced from a set* Γ *of assumptions*, as follows.

Definition 1.2.1 (Frege [1879]) *Given a set of axioms* $\mathcal{H}$*, the relation* $\vdash_{\mathcal{H}}$ *is inductively defined as follows:*

1. **Axioms**. *An axiom* $\beta \in \mathcal{H}$ *can be deduced from any set of assumptions:*
$$\Gamma \vdash_{\mathcal{H}} \beta.$$
2. **Assumptions**. *An assumption can be deduced from a set of assumptions to which it belongs:*
$$\Gamma, \beta \vdash_{\mathcal{H}} \beta.$$
3. **Modus Ponens**. *If* α *and* $\alpha \to \beta$ *are both deducible from a set of assumptions, then so is* β*:*
$$\frac{\Gamma \vdash_{\mathcal{H}} \alpha \qquad \Gamma \vdash_{\mathcal{H}} \alpha \to \beta}{\Gamma \vdash_{\mathcal{H}} \beta.}$$

Equivalence with Natural Deduction

The main difference between Hilbert systems and Natural Deduction is that the →-introduction rule of the latter (that allows the introduction of implications in general) is replaced by axioms that introduce particular implications, with the effect of turning an *analytical* (top-down) approach into a *synthetical* (bottom-up) one. Of course, the whole point is to get axioms that are sufficiently general to capture the essence of the →-introduction rule, in the sense of being able to reproduce it as a *derived rule*. The problem is that we have no clue as to which axioms will turn out to be sufficient. The best way is to attempt a proof of the equivalence

$$\Gamma \vdash_{\mathcal{H}} \beta \iff \Gamma \vdash_{\mathcal{N}} \beta,$$

and *discover* in the process which axioms are needed.

Suppose $\Gamma \vdash_{\mathcal{H}} \beta$. We proceed by induction on the definition 1.2.1. If $\beta \in \mathcal{H}$, i.e. β is an axiom, we will prove that $\vdash_{\mathcal{N}} \beta$; but first we have to find the axioms. If β is in Γ, then $\Gamma \vdash_{\mathcal{N}} \beta$ by definition 1.1.1. Finally, if β is obtained from $\alpha \to \beta$ and α by Modus Ponens, we can first of all apply the induction hypothesis, and suppose that we already have $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ and $\Gamma \vdash_{\mathcal{N}} \alpha$. Then an application of →-elimination produces $\Gamma \vdash_{\mathcal{N}} \beta$. In one word, if the axioms of a Hilbert system are provable in Natural Deduction then so are all its theorems, since the only deduction rule is Modus Ponens, and that is captured by →-elimination.

In the opposite direction, suppose $\Gamma \vdash_{\mathcal{N}} \beta$. We proceed by induction on the definition 1.1.1. If β is in Γ, then $\Gamma \vdash_{\mathcal{H}} \beta$ by definition 1.2.1. If $\Gamma \vdash_{\mathcal{N}} \beta$ is obtained by →-elimination, then (as above) $\Gamma \vdash_{\mathcal{H}} \beta$ by the induction hypothesis and Modus Ponens. If $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ is obtained by →-introduction, then the previous step was $\Gamma, \alpha \vdash_{\mathcal{N}} \beta$ and (by the induction hypothesis) $\Gamma, \alpha \vdash_{\mathcal{H}} \beta$. The proof of 1.2.3 will discover conditions on $\mathcal{H}$ under which $\Gamma \vdash_{\mathcal{H}} \alpha \to \beta$ follows (i.e., under which the →- introduction rule of Natural Deduction is a derived rule in a Hilbert system). Namely, it is enough that for any α, γ, δ the following are provable in $\mathcal{H}$:

1. $\alpha \to \alpha$
2. $\gamma \to (\alpha \to \gamma)$
3. $[(\alpha \to (\gamma \to \delta)] \to [(\alpha \to \gamma) \to (\alpha \to \delta)]$.

The easiest way to have these formulas provable, is to assume them as axioms.

We can thus state the result we were looking for.

Theorem 1.2.2 Equivalence of Hilbert Systems and Natural Deduction (Gentzen [1935]) *If $\mathcal{H}$ is any Hilbert system whose theorems include 1–3 of 1.2.3, then for any Γ and β:*

$$\Gamma \vdash_{\mathcal{H}} \beta \iff \Gamma \vdash_{\mathcal{N}} \beta.$$

Proof. 1 is provable in Natural Deduction as follows:

$$\frac{[\alpha]^{(1)}}{\alpha^{(1)} \to \alpha.}$$

2 is provable as follows:

$$\cfrac{\cfrac{[\alpha]^{(1)} \qquad [\beta]^{(2)}}{\alpha^{(1)} \to \beta}}{\beta^{(2)} \to (\alpha \to \beta).}$$

3 is provable as follows:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{[\alpha]^{(1)} \qquad [\alpha \to \beta]^{(2)}}{\beta} \qquad \cfrac{[\alpha]^{(1)} \qquad [\alpha \to (\beta \to \gamma)]^{(3)}}{\beta \to \gamma}}{\gamma}}{\alpha^{(1)} \to \gamma}}{(\alpha \to \beta)^{(2)} \to (\alpha \to \gamma)}}{[\alpha \to (\beta \to \gamma)]^{(3)} \to [(\alpha \to \beta) \to (\alpha \to \gamma)].}$$

This was the only gap left open in the discussion of the equivalence result. □

The Deduction Theorem

To complete the equivalence proof between Hilbert Systems and Natural Deduction we still need to prove the following announced result, whose proof will justify the choice of axioms for $\mathcal{H}$.

Theorem 1.2.3 Deduction Theorem (Herbrand [1928], Tarski [1930]) *If for any α, γ, δ the following are provable in $\mathcal{H}$:*

1. $\alpha \to \alpha$
2. $\gamma \to (\alpha \to \gamma)$
3. $[(\alpha \to (\gamma \to \delta)] \to [(\alpha \to \gamma) \to (\alpha \to \delta)]$,

then for any α, β and Γ:

$$\frac{\Gamma, \alpha \vdash_{\mathcal{H}} \beta}{\Gamma \vdash_{\mathcal{H}} \alpha \to \beta.}$$

Proof. By hypothesis, we have a proof of β using only Modus Ponens, and whose starting points are occurrences of either α, or axioms α_i, or assumptions $\gamma_j \in \Gamma$. The idea is to transform such a proof into a proof of $\alpha \to \beta$, by sticking '$\alpha \to$' in front of every formula in the original proof. We have three cases to consider.

The occurrences of α in the original proof become occurrences of $\alpha \to \alpha$, which is an instance of 1 above.

The occurrences of assumptions γ_j in the original proof become occurrences of $\alpha \to \gamma_j$, which can be deduced from the assumptions γ_j from instances of 2 above, by Modus Ponens:

$$\frac{\gamma_j \qquad \gamma_j \to (\alpha \to \gamma_j)}{\alpha \to \gamma_j.}$$

Similarly for the occurrences of axioms α_i.

The occurrences of applications of Modus Ponens

$$\frac{\gamma \qquad \gamma \to \delta}{\delta}$$

in the original proof become now occurrences of

$$\frac{\alpha \to \gamma \qquad \alpha \to (\gamma \to \delta)}{\alpha \to \delta.}$$

These are not anymore applications of Modus Ponens, but from $\alpha \to \gamma$ we can obtain $\alpha \to \delta$ by Modus Ponens from

$$(\alpha \to \gamma) \to (\alpha \to \delta),$$

and this can be obtained from $\alpha \to (\gamma \to \delta)$ by Modus Ponens from

$$[\alpha \to (\gamma \to \delta)] \to [(\alpha \to \gamma) \to (\alpha \to \delta)],$$

which is an instance of 3 above. In other words, we can expand

$$\frac{\alpha \to \gamma \qquad \alpha \to (\gamma \to \delta)}{\alpha \to \delta}$$

into the following segment of a proof by Modus Ponens:

$$\frac{\alpha \to \gamma \qquad \dfrac{\alpha \to (\gamma \to \delta) \qquad [\alpha \to (\gamma \to \delta)] \to [(\alpha \to \gamma) \to (\alpha \to \delta)]}{(\alpha \to \gamma) \to (\alpha \to \delta)}}{\alpha \to \delta.}$$

Then 1, 2 and 3 allow us to transform the original proof of β from the assumptions Γ and α into a proof of $\alpha \to \beta$ from the assumptions Γ. □

Exercises 1.2.4 Independence of the axioms.

a) *1 is derivable from 2 and 3*. (Hint: use the following axioms:

- $\alpha \to (\alpha \to \alpha)$

- $\alpha \to [(\alpha \to \alpha) \to \alpha]$
- $\{\alpha \to [(\alpha \to \alpha) \to \alpha]\} \to \{[\alpha \to (\alpha \to \alpha)] \to (\alpha \to \alpha)\}$.)

b) *None of 2 and 3 can be derived from the other*. (Hint: according to the third row of the following truth-table for implication, 3 always receives value T but 2 does not. According to the fourth row, 2 always receives value T but 3 does not.

α	T	T	T	F	F	F	U	U	U
β	T	F	U	T	F	U	T	F	U
$\alpha \to \beta$	T	U	U	U	U	T	T	T	T
$\alpha \to \beta$	T	F	U	T	T	F	T	T	T

Moreover, if α and $\alpha \to \beta$ always receive value T then so does β, in both cases.)

Exercise 1.2.5 *Transform the natural proof of $\alpha, \alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{H}} \gamma$, into a proof of $\alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{H}} \alpha \to \gamma$.)*

1.3 Sequents

Hilbert systems retained →-elimination, and presented an alternative approach to *→-introduction*, by substituting it with axioms from which all introductions could be synthesized. We now present an alternative approach to *→-elimination*, while retaining the rule of →-introduction.

The intuition for the new approach comes from an analysis of the following usual example:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{[\alpha]^{(1)} \quad [\alpha \to \beta]^{(3)}}{\beta} \qquad [\beta \to \gamma]^{(2)}}{\gamma}}{\alpha^{(1)} \to \gamma}}{(\beta \to \gamma)^{(2)} \to (\alpha \to \gamma)}}{(\alpha \to \beta)^{(3)} \to [(\beta \to \gamma) \to (\alpha \to \gamma)].}$$

Such a proof proceeds from top to bottom, by first analyzing a bunch of assumptions into atomic facts, and then recombining these facts into a compound conclusion.

The alternative approach proceeds from the center, i.e. from the atomic facts, and builds the proof by simultaneously working downward towards the conclusion, and upward towards the assumptions.

Consequence relation

The rules of the Sequent Calculus will only be →-introductions, but of two different kinds: in the conclusions and in the hypotheses. Inductively, having already a piece of a proof, we can either expand it at the bottom by complicating the conclusion, or at the top by complicating the assumptions. In the definition of $\vdash_{\mathcal{S}}$, this will correspond to introducing → on the right and on the left, respectively. This is a treatment of assumptions radically different from that of $\mathcal{N}$, where assumptions were fixed and could loose their status (by being discharged), but not be modified.

It is not surprising that complicating a conclusion corresponds to the usual →-introduction. Complicating an assumption β in a proof $\mathcal{D}_\gamma$ requires the introduction of a proof ending with β: this is naturally done by →-elimination. But if we only want to complicate β into $\alpha \to \beta$ without introducing the new assumption α, we will require a proof $\mathcal{D}_\alpha$ of α. With additional assumptions Γ, this is pictured as follows:

$$\begin{array}{lcc} & \Gamma & \\ & \mathcal{D}_\alpha & \\ & \alpha & \alpha \to \beta \\ \cline{2-3} \Gamma, & \multicolumn{2}{c}{\beta} \\ & \multicolumn{2}{c}{\mathcal{D}_\gamma} \\ & \multicolumn{2}{c}{\gamma.} \end{array}$$

We now give a formal inductive definition of the relation $\vdash_{\mathcal{S}}$.

Definition 1.3.1 (Gentzen [1935]) *The relation $\vdash_{\mathcal{S}}$ is inductively defined as follows:*

1. **Assumptions**. *An assumption can be deduced from a set of assumptions to which it belongs:*
$$\Gamma, \beta \vdash_{\mathcal{S}} \beta.$$

2. **→-Introduction on the right**. *If β is deducible from Γ and α, then $\alpha \to \beta$ is deducible from Γ:*
$$\frac{\Gamma, \alpha \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \to \beta.}$$

3. **→-Introduction on the left**. *If α is deducible from Γ and γ is deducible from Γ and β, then γ is deducible from Γ and $\alpha \to \beta$:*
$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \to \beta \vdash_{\mathcal{S}} \gamma.}$$

One nice feature of the rules of $\mathcal{S}$ is that they are *backward deterministic*, as opposed to →-elimination of $\mathcal{N}$ and Modus Ponens of $\mathcal{H}$ (that, when read backwards, introduce an arbitrary extraneous formula). Another way of expressing this is the **Subformula Property**: *in a proof of a sequent $\Gamma \vdash_{\mathcal{S}} \beta$, only subformulas of β and of formulas in Γ can occur*. A similar property does hold for $\mathcal{N}$ as well, but only for normal proofs (1.1.2.2).

The assumptions can be weakened to the following special cases:

$$\Gamma, p \vdash_{\mathcal{S}} p.$$

Then $\Gamma, \beta \vdash_{\mathcal{S}} \beta$ becomes derivable by induction on β, as follows. If $\beta = p$, then $\Gamma, p \vdash_{\mathcal{S}} p$ is given. And if $\beta = \alpha \to \gamma$, then

$$\dfrac{\dfrac{\Gamma, \alpha \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \alpha, \gamma \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha, \alpha \to \gamma \vdash_{\mathcal{S}} \gamma}}{\Gamma, \alpha \to \gamma \vdash_{\mathcal{S}} \alpha \to \gamma,}$$

where the sequents in the first line are given by the induction hypothesis.

The following **Thinning Rule** is a *derived rule*, in the sense that whenever we have a proof of the top sequent, we also have a proof of the bottom one (obtained from the previous one by inserting Δ on the left of every axiom, and continuing the proof as in the original one):

$$\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha}{\Gamma \cup \Delta \vdash_{\mathcal{S}} \alpha.}$$

Then *→-introduction on the left can be strengthened* as:

$$\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Delta, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma \cup \Delta, \alpha \to \beta \vdash_{\mathcal{S}} \gamma.}$$

This strengthened form will be tacitly used in the following, when convenient.

Equivalence with Natural Deduction

Since the axioms of $\mathcal{N}$ and $\mathcal{S}$ are the same, and the rule of →-introduction on the right of $\vdash_{\mathcal{S}}$ is the same as →-introduction of $\mathcal{N}$, the problem of the equivalence between the two systems is reduced to the provability of the rule of →-introduction on the left of $\mathcal{S}$ as a derived rule of $\mathcal{N}$, and of the →-elimination rule of $\mathcal{N}$ as a derived rule of $\mathcal{S}$.

There is no problem if we want to prove

$$\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma, \beta \vdash_{\mathcal{N}} \gamma}{\Gamma, \alpha \to \beta \vdash_{\mathcal{N}} \gamma.}$$

Indeed, by the hypothesis $\Gamma \vdash_{\mathcal{N}} \alpha$ we have a proof $\mathcal{D}_\alpha$ of α from the set of assumptions Γ. Similarly, we also have a proof $\mathcal{D}_\gamma$ of γ from the set of assumption Γ and β. We then first obtain a proof of β by →-elimination (using $\alpha \to \beta$ as an additional hypothesis), and then substitute it above every occurrence of β used in $\mathcal{D}_\gamma$:

$$\Gamma, \begin{array}{c} \begin{array}{cc} \begin{array}{c} \Gamma \\ \mathcal{D}_\alpha \\ \alpha \end{array} & \alpha \to \beta \end{array} \\ \hline \beta \\ \mathcal{D}_\gamma \\ \gamma. \end{array}$$

The assumptions of this proof are Γ and $\alpha \to \beta$, since the latter has not been discharged. We thus have $\Gamma, \alpha \to \beta \vdash_{\mathcal{N}} \gamma$.

There is instead a problem if we want to prove

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma \vdash_{\mathcal{S}} \alpha \to \beta}{\Gamma \vdash_{\mathcal{S}} \beta.}$$

The $\to$-introduction rule on the right allows us to introduce $\alpha \to \beta$, from the axiom $\Gamma, \beta \vdash_{\mathcal{S}} \beta$ and one of the hypotheses (namely, $\Gamma \vdash_{\mathcal{S}} \alpha$):

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \beta \vdash_{\mathcal{S}} \beta}{\Gamma, \alpha \to \beta \vdash_{\mathcal{S}} \beta.}$$

But we do not quite know what to do with this conclusion, i.e. how to match it with the remaining hypothesis $\Gamma \vdash_{\mathcal{S}} \alpha \to \beta$ to get the wanted conclusion $\Gamma \vdash_{\mathcal{S}} \beta$. Of course, there would be no problem if we had a rule that allowed us to cut a formula appearing in two sequents, once on the left and once on the right.

Definition 1.3.2 Cut Rule. *The system $\mathcal{S}+$ Cut is defined as the system $\mathcal{S}$, with the additional rule:*

$$\frac{\Gamma \vdash_{\mathcal{S}+\text{Cut}} \gamma \qquad \Gamma, \gamma \vdash_{\mathcal{S}+\text{Cut}} \beta}{\Gamma \vdash_{\mathcal{S}+\text{Cut}} \beta.}$$

Notice that the Cut Rule reintroduces the nondeterministic element that was present in $\mathcal{N}$ and $\mathcal{H}$, because of $\to$-elimination and Modus Ponens. Its justification is of course the fact that in $\mathcal{S}+$Cut we immediately get $\to$-elimination as a derived rule:

$$\frac{\Gamma \vdash_{\mathcal{S}+\text{Cut}} \alpha \to \beta \qquad \dfrac{\Gamma \vdash_{\mathcal{S}+\text{Cut}} \alpha \qquad \Gamma, \beta \vdash_{\mathcal{S}+\text{Cut}} \beta}{\Gamma, \alpha \to \beta \vdash_{\mathcal{S}+\text{Cut}} \beta}}{\Gamma \vdash_{\mathcal{S}+\text{Cut}} \beta.}$$

The equivalence of the two systems $\mathcal{N}$ and $\mathcal{S}$ would thus be proved if we had the following theorem.

Theorem 1.3.3 Cut Elimination (Gentzen [1935]) *For any Γ and β:*

$$\Gamma \vdash_{\mathcal{S}+\text{Cut}} \beta \implies \Gamma \vdash_{\mathcal{S}} \beta.$$

The intuitive reason why the Cut Elimination Theorem holds is *symmetry*. Notice that an axiom

$$\Gamma, \alpha \vdash \alpha$$

tells that an occurrence of α on the left of $\vdash$ implies an occurrence of α on the right. Symmetrically, the Cut Rule

$$\frac{\Gamma \vdash \alpha \qquad \Gamma, \alpha \vdash \beta}{\Gamma \vdash \beta}$$

tells that an occurrence of α on the right of $\vdash$ implies an occurrence on the left. By symmetry, *since the axioms are necessary, cuts are eliminable*.

The rest of the present section is devoted to giving two different proofs of this result, which provide different information. A third one will be given in 2.2.7. Any of these proofs will fill the remaining hole in the proof of the following result.

Theorem 1.3.4 Equivalence of the Sequent System and Natural Deduction (Gentzen [1935]) *For any* Γ *and* β*:*

$$\Gamma \vdash_{\mathcal{S}} \beta \iff \Gamma \vdash_{\mathcal{N}} \beta.$$

A Cut Elimination procedure

We first exhibit a procedure that shows how to operate directly on a proof of a given sequent in the system $\mathcal{S}$ + Cut, to transform it into a proof without cuts, by direct manipulations. The present proof of the Cut Elimination Theorem is thus constructive, and it also contains information on the complexity of the cut elimination procedure.

Theorem 1.3.5 Cut Elimination Procedure (Gentzen [1935]) *Every proof in the sequent system with cut can be transformed into a proof in the sequent system without cut, by means of an appropriate sequence of cut eliminations.*

Proof. We eliminate cuts one at a time by starting with uppermost ones, i.e. starting from cuts without any other cut above them. The reason is that then the part of the proof above the cut was done in the system $\mathcal{S}$ without Cut, and we are then able to go backwards in the proof, and know exactly what happened before the cut.

We show how to eliminate one such cut:

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \alpha \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \beta.}$$

The cut is *final* when one of the premises of the cut is an assumption. In these cases the cut is obviously redundant. Indeed, if $\Gamma \vdash_{\mathcal{S}} \alpha$ was introduced as an assumption then α belongs to Γ, and $\Gamma \vdash_{\mathcal{S}} \beta$ is equal to $\Gamma, \alpha \vdash_{\mathcal{S}} \beta$. If $\Gamma, \alpha \vdash_{\mathcal{S}} \beta$ was introduced as an assumption, then either $\beta = \alpha$, and then $\Gamma \vdash_{\mathcal{S}} \beta$ is equal to $\Gamma \vdash_{\mathcal{S}} \alpha$; or β belongs to Γ, and then $\Gamma \vdash_{\mathcal{S}} \beta$ can be introduced directly as an assumption.

The cut is *inductive* when α has just been introduced on both sides, that is: $\alpha = \gamma \to \delta$ for some γ and δ; $\Gamma \vdash_{\mathcal{S}} \gamma \to \delta$ has been obtained by $\to$-introduction on the right from $\Gamma, \gamma \vdash_{\mathcal{S}} \delta$; and $\Gamma, \gamma \to \delta \vdash_{\mathcal{S}} \beta$ has been obtained by $\to$-introduction on the left from $\Gamma \vdash_{\mathcal{S}} \gamma$ and $\Gamma, \delta \vdash_{\mathcal{S}} \beta$. Then the cut looks like this:

$$\dfrac{\dfrac{\Gamma, \gamma \vdash_{\mathcal{S}} \delta}{\Gamma \vdash_{\mathcal{S}} \gamma \to \delta} \qquad \dfrac{\Gamma \vdash_{\mathcal{S}} \gamma \qquad \Gamma, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \gamma \to \delta \vdash_{\mathcal{S}} \beta}}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \beta.}$$

The cut on $\gamma \to \delta$ can be eliminated as follows, by substituting it with *two* cuts, on the formulas γ and δ (of lower complexity):

$$\dfrac{\dfrac{\Gamma \vdash_{\mathcal{S}} \gamma \qquad \Gamma, \gamma \vdash_{\mathcal{S}} \delta}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \delta} \qquad \Gamma, \delta \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \beta.}$$

The appropriate measure of complexity is, as usual, the **degree** of a formula, defined as in 1.1.3.

The cut is *interlocutory* when one or both the occurrences of α have been introduced at steps preceeding the last ones (on the appropriate sides). In this case we can simply move the cut upwards, until it can be eliminated as above. In particular, we replace one cut on a given formula by *one* or *two* cuts on the same formula, but closer to the place of introduction of the latter. For example, a cut like

$$\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \dfrac{\Gamma, \alpha \vdash_{\mathcal{S}} \gamma \qquad \Gamma, \alpha, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \gamma \to \delta, \alpha \vdash_{\mathcal{S}} \beta}}{\Gamma, \gamma \to \delta \vdash_{\mathcal{S}+\mathrm{Cut}} \beta}$$

can be replaced by two as follows:

$$\dfrac{\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \alpha \vdash_{\mathcal{S}} \gamma}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \gamma} \qquad \dfrac{\Gamma, \delta \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \alpha, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \delta \vdash_{\mathcal{S}+\mathrm{Cut}} \beta}}{\Gamma, \gamma \to \delta \vdash_{\mathcal{S}} \beta,}$$

where $\Gamma, \delta \vdash_{\mathcal{S}} \alpha$ can be obtained from $\Gamma \vdash_{\mathcal{S}} \alpha$ by adding the assumption δ everywhere in the proof. The remaining cases are similar. □

Cut Elimination from Normalization

The proof of the Cut Elimination Theorem just given is obviously reminiscent of the proof of the Normalization Theorem given in Section 1.1. We are now going to show how we can actually translate cut-free proofs of $\mathcal{S}$ into normal proofs of $\mathcal{N}$ and conversely. Thus *the Cut Elimination Theorem for $\mathcal{S}$ and the Normalization Theorem for $\mathcal{N}$ are equivalent results.* For example, to prove the former from the latter first note that $\mathcal{N}$ is closed under cut: indeed,

$$\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma, \alpha \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \beta}$$

corresponds to the step

$$\text{from}\qquad \begin{array}{c}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{array}\qquad\text{and}\qquad \begin{array}{c}\Gamma,\alpha\\ \mathcal{D}_\beta\\ \beta\end{array}\qquad\text{to}\qquad \begin{array}{c}\Gamma\\ \mathcal{D}_\alpha\\ \Gamma,\quad\alpha\\ \mathcal{D}_\beta\\ \beta.\end{array}$$

Then proofs in the system $\mathcal{S} + \text{Cut}$ can be translated into proofs of $\mathcal{N}$, normalized and retranslated back into cut-free proofs of $\mathcal{S}$.

The precise result is the following.

Proposition 1.3.6 (Prawitz [1965]) *There are canonical translations of cut-free proofs in $\mathcal{S}$ to normal proofs in $\mathcal{N}$, and conversely.*

Proof. The translation from $\mathcal{S}$ to $\mathcal{N}$ given as half of the proof of the equivalence of the two systems (1.3.4) already shows that (cut-free) proofs in $\mathcal{S}$ correspond to normal proofs in $\mathcal{N}$.

For the converse, the translation from $\mathcal{N}$ to sequents given in 1.3.4 is not useful here, since it uses the cut rule (which was introduced precisely for the purpose of that translation). That translation proceeds from top down, i.e. by forward induction on the construction of the proof (starting from the first steps). We now proceed from bottom up, i.e. by backward induction on the construction of the proof (starting from the last step).

Suppose we have a normal proof of β from Γ in $\mathcal{N}$. We proceed inductively on the length of the proof of $\Gamma \vdash_{\mathcal{N}} \beta$ and show that $\Gamma \vdash_{\mathcal{S}} \beta$, i.e. β can be deduced from Γ in $\mathcal{S}$ by a cut-free proof. There are three cases:

1. If $\Gamma \vdash_{\mathcal{N}} \beta$ is an assumption then β belongs to Γ, and thus $\Gamma \vdash_{\mathcal{S}} \beta$.

2. If $\beta = \gamma \to \delta$ and $\Gamma \vdash_{\mathcal{N}} \gamma \to \delta$ has been obtained by →-introduction from a normal proof of $\Gamma, \gamma \vdash_{\mathcal{N}} \delta$, then $\Gamma, \gamma \vdash_{\mathcal{S}} \delta$ by the induction hypothesis, and thus $\Gamma \vdash_{\mathcal{S}} \gamma \to \delta$.

3. If $\Gamma \vdash_{\mathcal{N}} \beta$ has been obtained by →-elimination from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ by a normal proof, then $\alpha \to \beta$ cannot have been obtained by →-introduction in the last step of its proof, otherwise it would be a maximum (because the next step is a →-elimination). Then either $\alpha \to \beta$ is itself an assumption, or it is obtained by →-elimination. By continuing upwards in the given normal proof, we eventually reach an assumption $\gamma \to \delta$ (see 1.1.2.1). We can thus

suppose that $\Gamma \vdash_{\mathcal{N}} \beta$ has the following form:

$$\Gamma,\ \dfrac{\begin{array}{c}\Gamma\\ \mathcal{D}_\gamma\\ \gamma\end{array} \qquad \gamma \to \delta}{\begin{array}{c}\delta\\ \mathcal{D}_\beta\\ \beta.\end{array}}$$

Notice that $\mathcal{D}_\gamma$ depends only on Γ. Indeed, other assumptions would have to be discharged later on, but this is impossible because there are no introduction rules and no minor premises in the path from $\gamma \to \delta$ to β.

We can now apply the induction hypothesis to both $\mathcal{D}_\gamma$ and $\mathcal{D}_\beta$, and get cut-free proofs $\Gamma \vdash_{\mathcal{S}} \gamma$ and $\Gamma, \delta \vdash_{\mathcal{S}} \beta$. By →-introduction on the left,

$$\dfrac{\Gamma \vdash_{\mathcal{S}} \gamma \qquad \Gamma, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \gamma \to \delta \vdash_{\mathcal{S}} \beta.}$$

But $\gamma \to \delta$, being an assumption, is already in Γ. The conclusion is thus equivalent to $\Gamma \vdash_{\mathcal{S}} \beta$. □

To give an example of the translation just introduced, consider the following normal proof:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{[\alpha]^{(1)} \quad [\alpha\to\beta]^{(2)}}{\beta} \qquad \dfrac{[\alpha]^{(1)} \quad [\alpha\to(\beta\to\gamma)]^{(3)}}{\beta\to\gamma}}{\gamma}}{\alpha^{(1)}\to\gamma}}{(\alpha\to\beta)^{(2)}\to(\alpha\to\gamma)}}{(\alpha\to(\beta\to\gamma))^{(3)}\to((\alpha\to\beta)\to(\alpha\to\gamma)).}$$

The last three steps are →-introductions for $\mathcal{N}$, which are the same as →-introductions on the right for $\mathcal{S}$. We thus only have to worry about the →-eliminations, which are treated inductively as follows. First,

$$\dfrac{\dfrac{\alpha \quad \alpha\to\beta}{\beta} \qquad \dfrac{\alpha \quad \alpha\to(\beta\to\gamma)}{\beta\to\gamma}}{\gamma}$$

is translated as:

$$\dfrac{\alpha \vdash_{\mathcal{S}} \alpha \qquad \alpha, \alpha\to\beta, \beta\to\gamma \vdash_{\mathcal{S}} \gamma}{\alpha, \alpha\to\beta, \alpha\to(\beta\to\gamma) \vdash_{\mathcal{S}} \gamma.}$$

Then

$$\dfrac{\dfrac{\alpha \quad \alpha\to\beta}{\beta} \qquad \beta\to\gamma}{\gamma}$$

is translated as:

$$\dfrac{\alpha, \alpha \to \beta \vdash_{\mathcal{S}} \beta \qquad \gamma \vdash_{\mathcal{S}} \gamma}{\alpha, \alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{S}} \gamma.}$$

Finally,

$$\dfrac{\alpha \qquad \alpha \to \beta}{\beta}$$

is translated as:

$$\dfrac{\alpha \vdash_{\mathcal{S}} \alpha \qquad \beta \vdash_{\mathcal{S}} \beta}{\alpha, \alpha \to \beta \vdash_{\mathcal{S}} \beta.}$$

By putting everything together, we get the following translation of the original proof:

$$\dfrac{\dfrac{\dfrac{\dfrac{\alpha \vdash_{\mathcal{S}} \alpha \qquad \dfrac{\dfrac{\alpha \vdash_{\mathcal{S}} \alpha \qquad \beta \vdash_{\mathcal{S}} \beta}{\alpha, \alpha \to \beta \vdash_{\mathcal{S}} \beta} \qquad \gamma \vdash_{\mathcal{S}} \gamma}{\alpha, \alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{S}} \gamma}}{\alpha, \alpha \to \beta, \alpha \to (\beta \to \gamma) \vdash_{\mathcal{S}} \gamma}}{\alpha \to \beta, \alpha \to (\beta \to \gamma) \vdash_{\mathcal{S}} \alpha \to \gamma}}{\alpha \to (\beta \to \gamma) \vdash_{\mathcal{S}} (\alpha \to \beta) \to (\alpha \to \gamma)}}{\vdash_{\mathcal{S}} (\alpha \to (\beta \to \gamma)) \to ((\alpha \to \beta) \to (\alpha \to \gamma)).}$$

It should be noted that the two translations from cut-free proofs to normal proofs and back are not (and cannot be) inverse of each other, since *the translation from $\mathcal{S}$ to $\mathcal{N}$ is not one-one*: the same normal proof can be the translation of different cut-free proofs. The reason is that sequent proofs specify an *order* in the construction of the corresponding normal proof, but no trace of this order remains after the proof has been constructed. For example, the normal proof

$$\dfrac{\dfrac{\alpha \qquad \alpha \to \beta}{\beta} \qquad \beta \to \gamma}{\gamma}$$

is the translation of the following two proofs of the sequent

$$\alpha, \alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{S}} \gamma,$$

originated by two different orders of →-introduction on the left:

$$\dfrac{\dfrac{\alpha \vdash_{\mathcal{S}} \alpha \qquad \beta \vdash_{\mathcal{S}} \beta}{\alpha, \alpha \to \beta \vdash_{\mathcal{S}} \beta} \qquad \gamma \vdash_{\mathcal{S}} \gamma}{\alpha, \alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{S}} \gamma,}$$

and

$$\dfrac{\alpha \vdash_{\mathcal{S}} \alpha \qquad \dfrac{\beta \vdash_{\mathcal{S}} \beta \qquad \gamma \vdash_{\mathcal{S}} \gamma}{\beta, \beta \to \gamma \vdash_{\mathcal{S}} \gamma}}{\alpha, \alpha \to \beta, \beta \to \gamma \vdash_{\mathcal{S}} \gamma.}$$

æ

Chapter 2

Semantics

The three approaches to implication introduced in Chapter 1 were syntactical, and captured different aspects of its constructive meaning. We turn now to a semantical characterization. In Section 1 we introduce the classical approach, and show its limitations. In Sections 2 and 3 we modify the classical approach into a more appropriate one.

2.1 Classical Semantics

Intuitively the implication $\alpha \to \beta$ means that whenever we know α, then we know β. This obviously imposes a restriction: under the intuitive notion of truth, if both α and $\alpha \to \beta$ are true, then so must be β. Equivalently, if α is true and β is false, then $\alpha \to \beta$ cannot be true. This necessary condition is turned into a necessary and sufficient one by classical logic, according to which $\alpha \to \beta$ is true in all other cases.

Definition of truth

The classical approach to truth is simple. We define *contingent* worlds as possible truth-value configurations of the linguistic atoms (the propositional letters), and determine the contingent truth or falsity of every formula. Then *necessary* or *absolute* truth (validity) is defined as truth in every possible worlds.

More formally, a **classical possible world** is a subset $\mathcal{A}$ of the propositional letters and it determines, by induction, the truth-value of any formula α, as in the following definition.

Definition 2.1.1 Classical Truth (Tarski [1936]) *The relation $\models$ is inductively defined as follows.*

$$\begin{array}{lcl} \mathcal{A} \models p & \Leftrightarrow & p \in \mathcal{A} \\ \mathcal{A} \models \alpha \to \beta & \Leftrightarrow & (\mathcal{A} \models \alpha \ \Rightarrow \ \mathcal{A} \models \beta). \end{array}$$

We read $\mathcal{A} \models \alpha$ as 'α is **true** *in $\mathcal{A}$', or '$\mathcal{A}$ is a* **model** *of α'. We also write $\mathcal{A} \not\models \alpha$ for the negation of $\mathcal{A} \models \alpha$, and we read it as '$\alpha$ is* **false** *in $\mathcal{A}$'.*

Note that the implication $\Rightarrow$ used in the definition is intended to be false only in the case when the premise is true but the conclusion is false, and true in all other cases.

There is no circularity in using $\Rightarrow$ to define the meaning of $\to$: both of them are implications, but the former is used *informally* (in technical terms: metalinguistically), with an intended meaning, to talk about the *formal* one (used in the language). The definition of $\models$ is meant to force the meaning of $\to$ to mirror the meaning of $\Rightarrow$, at the appropriate level.

We turn now to global truth, independent of the world.

Definition 2.1.2 *A formula α is a* **logical consequence** *of Γ (written $\boldsymbol{\Gamma \models \alpha}$) if α is true in every world in which all formulas of Γ are true.*

In the limit case of Γ empty, we get the notion of validity: α is **valid** (written $\models \alpha$) if α is true in every world.

Soundness

The next result shows that classical validity is an upper bound to provability in any of the equivalent systems for intuitionistic logic considered so far.

Theorem 2.1.3 Classical Soundness. *For any Γ and α:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \ \Rightarrow \ \Gamma \models \alpha.$$

Proof. By induction on the definition 1.1.1.

If $\Gamma, \beta \vdash_{\mathcal{N}} \beta$ is an assumption, then $\Gamma, \beta \models \beta$ is trivially true.

If $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ is obtained from $\Gamma, \alpha \vdash_{\mathcal{N}} \beta$ by $\to$-introduction, then $\Gamma, \alpha \models \beta$ by the induction hypothesis. Let $\mathcal{A}$ be any world that makes all formulas of Γ true. If $\mathcal{A}$ makes α false, then it makes $\alpha \to \beta$ true by definition 2.1.1. By the induction hypothesis, if $\mathcal{A}$ makes α true, then it must make β true, and hence $\alpha \to \beta$ true. Thus $\Gamma \models \alpha \to \beta$.

If $\Gamma \vdash_{\mathcal{N}} \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ by $\to$-elimination, then $\Gamma \models \alpha$ and $\Gamma \models \alpha \to \beta$ by the induction hypothesis. Let $\mathcal{A}$ be any world that makes all formulas of Γ true. Then $\mathcal{A}$ makes α and $\alpha \to \beta$ true by the induction hypothesis, and hence β true by definition 2.1.1. Thus $\mathcal{A} \models \beta$. □

A counterexample to completeness

The previous theorem shows that the equivalent provability notions introduced in this chapter are sound for classical semantics. The next formula, called *Peirce's Law*,[1] shows that the converse fails:

$$[(p \to q) \to p] \to p. \tag{2.1}$$

Indeed, consider any world $\mathcal{A}$: p is either true or false in it. In the first case 2.1 is true (an implication with true consequence). In the second case: $p \to q$ is true (false premise), $(p \to q) \to p$ is false (true premise and false conclusion) and 2.1 is true (false premise). Thus 2.1 is valid.

But 2.1 is not provable in $\mathcal{N}$. This is immediate using the Sequent System and the Cut Elimination Theorem 1.3.3. Indeed, the only possible cut-free proof of 2.1 in $\mathcal{S}$ would be

$$\dfrac{\dfrac{\dfrac{p \vdash_{\mathcal{S}} q}{\vdash_{\mathcal{S}} p \to q} \qquad p \vdash_{\mathcal{S}} p}{(p \to q) \to p \vdash_{\mathcal{S}} p}}{\vdash_{\mathcal{S}} [(p \to q) \to p] \to p.}$$

However, this is not a proof because $p \vdash_{\mathcal{S}} q$ is not an axiom.

2.2 Beth-Kripke Semantics

The example of Peirce's Law shows that the classical notion of truth is too simpleminded to capture the essence of implication, as the latter is defined by any of the equivalent systems $\mathcal{N}$, $\mathcal{H}$ and $\mathcal{S}$. The intuitive reason is easy to formulate, by looking once again at the intended meaning of implication: $\alpha \to \beta$ expresses the fact that *whenever* we know α, then we know β. In classical worlds, there is no 'whenever': for what concerns knowledge, there is only one instant, in which any formula is either true or false. In particular, there is no need for implication itself: we either know α, and hence also β, or we don't, and then the implication $\alpha \to \beta$ cannot be used.

Classical possible worlds are thus worlds for omniscient gods, whose knowledge is complete and never changes. Implication is instead a human concept: we prove $\alpha \to \beta$ as a step towards β, that will be completed if and when we get to know α. To capture the idea of changing *states of knowledge*, we introduce possible worlds in which knowledge can expand.

The new semantics will thus be *temporal* and *epistemic*, in the sense that it will deal with the concept of knowing something at a given time.

[1] Peirce's Law will be put into a broader context in 21.2.4.

Beth-Kripke models

The first assumption we make is that *knowledge is monotone*, in the sense that we never forget what we already know. We are thus modelling social, rather than personal knowledge: not what somebody knows and can forget, but what 'is known' at a given instant. Of course, there are different and *incompatible ways of extending knowledge* at a given instant: even when building on a fixed past experience, there are different directions research can take. A tree of possibilities will be a possible world, and a branch of such a tree will determine a possible history of knowledge in that particular world.

Recall that classical knowledge was determined by a set $\mathcal{A}$ of propositional letters, representing the true atomic facts and determining, by exclusion, the false ones as well. Instantaneous knowledge will now still be represented by sets $\mathcal{A}_\sigma$ of propositional letters, representing the true atomic facts known in a given knowledge state σ. But since knowledge can be expanded, letters not in $\mathcal{A}_\sigma$ are better thought of as unknown at that instant, rather than false.

Definition 2.2.1 (Beth [1956], Kripke [1963]) *An* **intuitionistic possible world***, is a triple*

$$\mathcal{A} = \langle P_\mathcal{A}, \sqsubseteq_\mathcal{A}, \{\mathcal{A}_\sigma\}_{\sigma \in P_\mathcal{A}} \rangle$$

where:

1. *$P_\mathcal{A}$ is a nonempty set of elements, representing states of knowledge*
2. *$(P_\mathcal{A}, \sqsubseteq_\mathcal{A})$ is a partial ordering*
3. *for each $\sigma \in P_\mathcal{A}$, $\mathcal{A}_\sigma$ is a set of propositional letters, representing the atomic facts known in state σ*
4. *$\sigma \sqsubseteq_\mathcal{A} \tau \;\Rightarrow\; \mathcal{A}_\sigma \subseteq \mathcal{A}_\tau$.*

Forcing

We now rephrase the definition of global truth in a given classical world (2.1.1) into a definition of local truth, relative also to given knowledge states

Definition 2.2.2 Forcing (Cohen [1963], Kripke [1963]) *For a given possible world $\mathcal{A}$, the relation $\Vdash_\mathcal{A}$ is inductively defined as follows.*

$$\begin{array}{lcl} \sigma \Vdash_\mathcal{A} p & \Leftrightarrow & p \in \mathcal{A}_\sigma \\ \sigma \Vdash_\mathcal{A} \alpha \to \beta & \Leftrightarrow & (\forall \tau \sqsupseteq_\mathcal{A} \sigma)(\tau \Vdash_\mathcal{A} \alpha \;\Rightarrow\; \tau \Vdash_\mathcal{A} \beta). \end{array}$$

We read $\sigma \Vdash_\mathcal{A} \alpha$ as 'α is **true** *in $\mathcal{A}$ at σ', or 'σ* **forces** *α in $\mathcal{A}$'.*

Forcing gives $\alpha \to \beta$ a meaning closer to the intended one. More precisely, $\alpha \to \beta$ becomes true as soon as we come to know that, whenever in the future we will know α, then we will know β.

Notice that the definition of forcing for intuitionistic implication involves a classical quantifier and a classical implication: with forcing we thus interpret *intuitionistic propositional* logic by means of *classical first order* logic, in particular in a very non constructive way. Thus the proposed semantics captures only certain aspects of the intended meaning of intuitionistic implication.

Classical truth is obviously a special case of forcing, in which $P_{\mathcal{A}}$ consists of just one element a, and $\mathcal{A}_a$ is the final knowledge. Then

$$\mathcal{A} \models \alpha \Leftrightarrow a \Vdash_{\mathcal{A}} \alpha,$$

because the definition of forcing for implication reduces, inductively, to

$$a \Vdash_{\mathcal{A}} \alpha \to \beta \Leftrightarrow (\mathcal{A} \models \alpha \Rightarrow \mathcal{A} \models \beta).$$

The following is a trivial but crucial property.

Proposition 2.2.3 Monotonicity of Forcing. *If a formula is forced at a given state, it remains forced at every following state:*

$$\sigma \Vdash_{\mathcal{A}} \alpha \wedge \sigma \sqsubseteq_{\mathcal{A}} \tau \Rightarrow \tau \Vdash_{\mathcal{A}} \alpha.$$

Proof. By induction on α. The atomic case holds by monotonicity of knowledge. The case of an implication follows from transitivity of $\sqsubseteq_{\mathcal{A}}$. □

Intuitionistic validity

We now turn to global truth, independently of the world and of the states of knowledge in it.

For shortness, we say that a formula α is **forced in a world $\mathcal{A}$** if it is forced at every state, i.e. if $(\forall\sigma \in P_{\mathcal{A}})(\sigma \Vdash \alpha)$.

Definition 2.2.4 *A formula α is an* **intuitionistic logical consequence** *of Γ (written $\boldsymbol{\Gamma \models_i \alpha}$) if α is forced in every world in which all formulas of Γ are forced.*

A formula α is **intuitionistically valid** *(written $\boldsymbol{\models_i \alpha}$) if α is forced in every world.*

Formally, the previous definition amounts to following:

$$(\forall\mathcal{A})[(\forall\sigma \in P_{\mathcal{A}})(\sigma \Vdash \Gamma) \implies (\forall\sigma \in P_{\mathcal{A}})(\sigma \Vdash \alpha),$$

where $\sigma \Vdash \Gamma$ means $(\forall\gamma \in \Gamma)(\sigma \Vdash \gamma)$.

An alternative, apparently stronger definition would be the following:

$$(\forall \mathcal{A})(\forall \sigma \in P_{\mathcal{A}})(\sigma \Vdash \Gamma \implies \sigma \Vdash \alpha),$$

i.e. α is forced at every state in which all formulas of Γ are forced.

Using the stronger definition produces a stronger Soundness Theorem, and using the weaker one produces a stronger Completeness Theorem. This is what we do in the following two proofs, thus proving in passing an equivalence of the two definitions.

Soundness and Completeness

Since classical worlds are special cases of intuitionistic ones, the next result extends the Classical Soundness Theorem, and it is proved in a similar way.

Theorem 2.2.5 Intuitionistic Soundness (Beth [1956], Kripke [1963]) *For any* Γ *and* α*:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \ \Rightarrow\ \Gamma \models_i \alpha.$$

Proof. By induction on the definition 1.1.1.

If $\Gamma, \beta \vdash_{\mathcal{N}} \beta$ is an assumption, then $\Gamma, \beta \models_i \beta$ is trivially true.

If $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ is obtained from $\Gamma, \alpha \vdash_{\mathcal{N}} \beta$ by $\to$-introduction, then $\Gamma, \alpha \models_i \beta$ by the induction hypothesis. Let σ be any state that forces all formulas of Γ in some world $\mathcal{A}$, and let τ be any extension of σ. If τ does not force α, then the implication

$$\tau \Vdash_{\mathcal{A}} \alpha \ \Rightarrow\ \tau \Vdash_{\mathcal{A}} \beta$$

is true. If τ forces α then, since τ also forces all formulas of Γ by monotonicity (because it extends σ), τ forces β by the induction hypothesis. Thus σ forces $\alpha \to \beta$ by definition of forcing. Since σ and $\mathcal{A}$ are arbitrary, $\Gamma \models_i \alpha \to \beta$.

If $\Gamma \vdash_{\mathcal{N}} \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ by $\to$-elimination, then $\Gamma \models_i \alpha$ and $\Gamma \models_i \alpha \to \beta$ by the induction hypothesis. Let σ be any state that forces all formulas of Γ in some world $\mathcal{A}$. By the induction hypotheses, σ forces α and $\alpha \to \beta$. By definition of forcing, then σ forces β. Since σ and $\mathcal{A}$ are arbitrary, $\Gamma \models_i \beta$. □

The next theorem tells us that the intuition that led us to the extended notion of intuitionistic world was correct.

Theorem 2.2.6 Intuitionistic Completeness (Beth [1956], Kripke [1963]) *For any* Γ *and* α*:*

$$\Gamma \models_i \alpha \ \Rightarrow\ \Gamma \vdash_{\mathcal{N}} \alpha.$$

Proof. We build a single possible world $\mathcal{A}$ such that if $\Gamma \not\vdash_{\mathcal{N}} \alpha$, then there is a state Θ in it such that Θ forces all formulas in Γ, but not α. Then the restriction of $\mathcal{A}$ to all states above Θ provides a world in which all formulas in Γ are forced, but α is not. This proves the contrapositive of the stated result.

The idea is to consider the sets of consequences of any possible Γ as the states of knowledge. Thus the world $\mathcal{A}$ is defined as follows:

$$\mathcal{A} = \langle \mathcal{F}, \subseteq, \{\mathcal{A}_\Theta\}_{\Theta \in \mathcal{F}} \rangle,$$

where:

1. $\mathcal{F}$ is the set of all sets of formulas Θ closed under $\vdash_{\mathcal{N}}$, i.e. such that if $\Gamma \subseteq \Theta$ and $\Gamma \vdash_{\mathcal{N}} \alpha$, then $\alpha \in \Theta$

2. $\subseteq$ is the usual set-theoretical inclusion relation

3. $\mathcal{A}_\Theta$ is the set of formulas in Θ consisting only of a propositional letter, i.e.

$$\mathcal{A}_\Theta = \{p : p \in \Theta\}.$$

The main point is that, due to closure under $\vdash_{\mathcal{N}}$, *forcing is reduced to membership*: for any formula α,

$$\Theta \Vdash_{\mathcal{A}} \alpha \ \Leftrightarrow \ \alpha \in \Theta.$$

For $\alpha = p$, this holds by the definition of forcing and of $\mathcal{A}_\Theta$. For $\alpha = \gamma \to \delta$ it is proved inductively as follows, using $\to$-introduction in one direction and $\to$-elimination in the other:

- Suppose $\Theta \Vdash_{\mathcal{A}} \gamma \to \delta$, i.e.

$$(\forall \Delta \supseteq \Theta)(\Delta \Vdash_{\mathcal{A}} \gamma \ \Rightarrow \ \Delta \Vdash_{\mathcal{A}} \delta).$$

By the induction hypothesis,

$$(\forall \Delta \supseteq \Theta)(\gamma \in \Delta \ \Rightarrow \ \delta \in \Delta).$$

We want $\gamma \to \delta \in \Theta$. Suppose $\gamma \to \delta \notin \Theta$. By closure under $\vdash_{\mathcal{N}}$, $\Theta \not\vdash_{\mathcal{N}} \gamma \to \delta$. By $\to$*-introduction*, $\Theta, \gamma \not\vdash_{\mathcal{N}} \delta$. Let Δ be the closure under $\vdash_{\mathcal{N}}$ of $\Theta \cup \{\gamma\}$. Then $\Delta \supseteq \Theta$ and thus, by the induction hypothesis, if $\gamma \in \Delta$ then $\delta \in \Delta$. But $\gamma \in \Delta$ (by definition), while $\delta \notin \Delta$ (because Δ is the closure under $\vdash_{\mathcal{N}}$ of $\Theta \cup \{\gamma\}$, and $\Theta, \gamma \not\vdash_{\mathcal{N}} \delta$), contradiction.

- Suppose $\gamma \to \delta \in \Theta$. We want $\Theta \Vdash_{\mathcal{A}} \gamma \to \delta$, i.e.

$$(\forall \Delta \supseteq \Theta)(\Delta \Vdash_{\mathcal{A}} \gamma \ \Rightarrow \ \Delta \Vdash_{\mathcal{A}} \delta)$$

or, by the induction hypothesis,

$$(\forall \Delta \supseteq \Theta)(\gamma \in \Delta \;\Rightarrow\; \delta \in \Delta).$$

Let $\Delta \supseteq \Theta$, in particular $\gamma \to \delta \in \Delta$. If $\gamma \in \Delta$, then $\delta \in \Delta$ by $\to$-*elimination* and closure under $\vdash_{\mathcal{N}}$.

Suppose now that $\Gamma \not\vdash_{\mathcal{N}} \alpha$, and let Θ be the closure of Γ under $\vdash_{\mathcal{N}}$. We have proved that

$$\Theta \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \alpha \in \Theta.$$

Since $\alpha \notin \Theta$ (because $\Gamma \not\vdash_{\mathcal{N}} \alpha$), it follows that Θ does not force α. Since Θ does force every formula in Γ, being an extension of it, it follows that $\Gamma \not\models_i \alpha$. □

Notice that *the proof of the Intuitionistic Completeness Theorem is constructive*. This is not automatic from the proof given above, since we actually proved the contrapositive of the needed statement. But we will prove in 3.1.1 that $\mathcal{N}$ is decidable. Then, given Γ and α, we can first of all decide whether $\Gamma \vdash_{\mathcal{N}} \alpha$ holds. If so, $\Gamma \models_i \alpha$ by the Intuitionistic Soundedness Theorem. Otherwise, the proof just given shows that $\Gamma \not\models_i \alpha$. This proves in a constructive way that if $\Gamma \vdash_{\mathcal{N}} \alpha$, then $\Gamma \models_i \alpha$.

Notice also that the proof provides *a single world* such that if α is not provable in $\mathcal{N}$, then it is not forced in it. Unfortunately, this single world is pretty complicated because it has uncountably many states (all possible sets of formulas closed under $\vdash_{\mathcal{N}}$). We will present improvements of this result in a short while.

A possible use of the Intuitionistic Completeness Theorem is to show that a formula is not intuitionistically provable, by exhibiting an intuitionistic world and a state in which the formula is not forced. For example, to show that Peirce's Law is not provable, it is enough to consider a world $\mathcal{A}$ with two states $\emptyset$ and 0, and such that $\mathcal{A}_{\emptyset} = \emptyset$ and $\mathcal{A}_0 = \{p\}$. No state forces $p \to q$, because each state has an extension (namely, 0) forcing p but not q. So every state forces $(p \to q) \to p$. Then $\emptyset$ does not force $[(p \to q) \to p] \to p$, because it forces $(p \to q) \to p$ but not p.

A semantical proof of Cut Elimination

From the Intuitionistic Completeness Theorem we immediately get a *semantical proof of the Cut Elimination Theorem.* We only have to check that the Cut Rule is a derived rule for $\models_i$. By the Intuitionistic Soundness Theorem, this will imply that if $\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \alpha$, then $\Gamma \models_i \alpha$. By the Intuitionistic Completeness Theorem and the equivalence of $\mathcal{N}$ and $\mathcal{S}$, we will then have $\Gamma \vdash_{\mathcal{S}} \alpha$.

Proposition 2.2.7 *Cut is a derived rule for* $\models_i$, *i.e.*

$$\frac{\Gamma \models_i \alpha \qquad \Gamma, \alpha \models_i \beta}{\Gamma \models_i \beta.}$$

Proof. Suppose $\Gamma \models_i \alpha$ and $\Gamma, \alpha \models_i \beta$. By the former, α must be forced at any σ forcing all formulas of Γ in any $\mathcal{A}$. By the latter, then β must be forced at any such σ. Hence $\Gamma \models_i \beta$ by definition. □

Notice that, as it stands, this proof is not very interesting: it uses the equivalence of $\mathcal{N}$ and $\mathcal{S}$, whose proof already required the Cut Elimination Theorem! But we can prove the Intuitionistic Completeness Theorem directly for $\mathcal{S}$, with the same proof as in 2.2.6.

Even after the adjustment just described, the present proof of Cut Elimination remains quite indirect: given a proof of a sequent in the system with Cut, it only tells that a proof in the system without Cut exists, but it does not show how to obtain it (except for the trivial and inefficient method of generating all possible proofs, until one is found).

In more complicated contexts, e.g. in Second Order Logic, this indirect feature is balanced by an advantage: an analogue of the less informative semantical proof is substantially easier to obtain than an analogue of the more informative, but more complicated syntactical proof in the style of Section 1.3.

Refinements of the Completeness Theorem ⋆

We now consider restrictions on the underlying partial orderings of intuitionistic worlds, to see whether they still provide a class of worlds sufficiently rich to determine intuitionistic completeness.

Exercises 2.2.8 a) *Trees are enough.* (Hint: replace a world

$$\mathcal{A} = \langle P_{\mathcal{A}}, \sqsubseteq_{\mathcal{A}}, \{\mathcal{A}_\sigma\}_{\sigma \in P_{\mathcal{A}}} \rangle$$

with $(P_{\mathcal{A}}, \sqsubseteq_{\mathcal{A}})$ a partial ordering, by a world

$$\mathcal{T} = \langle T_{\mathcal{T}}, \sqsubseteq_{\mathcal{T}}, \{\mathcal{T}_\sigma\}_{\sigma \in P_{\mathcal{T}}} \rangle$$

where: $T_{\mathcal{T}}$ is the set of all possible finite sequences $\langle \sigma_1, \ldots, \sigma_n \rangle$ of distinct elements of $P_{\mathcal{A}}$ such that $\sigma_1 \sqsubseteq_{\mathcal{A}} \sigma_2 \sqsubseteq_{\mathcal{A}} \cdots \sqsubseteq_{\mathcal{A}} \sigma_n$; $\sqsubseteq_{\mathcal{T}}$ is the order of sequences by extension; and $\mathcal{T}_{\langle \sigma_1, \ldots, \sigma_n \rangle} = \mathcal{A}_{\sigma_n}$.)

b) *Linear orderings are not enough.* (Hint: this uses results proved later, as follows. By 5.3.2.d, any linear Kripke model forces $(\alpha \to \beta) \vee (\beta \to \alpha)$, which is not intuitionistically valid. Thus linear Kripke models are not enough for completeness.

To translate the disjunction into an implicational formula, first note that it is classically valid, and hence its double negation is intuitionistically valid by 21.1.3. By a series of intuitionistically valid equivalences, in particular the ‘Good’ De Morgan Law discussed in Section 21.2, we get:

$$\begin{aligned} \neg\neg[(\alpha \to \beta) \vee (\beta \to \alpha)] &\Leftrightarrow \neg[\neg(\alpha \to \beta) \wedge \neg(\beta \to \alpha)] \\ &\Leftrightarrow \{[(\alpha \to \beta) \to \bot] \wedge [(\beta \to \alpha) \to \bot]\} \to \bot \\ &\Leftrightarrow [(\alpha \to \beta) \to \bot) \to \{[(\beta \to \alpha) \to \bot] \to \bot\}. \end{aligned}$$

By replacing $\bot$ by any formula γ, we get a stronger formula

$$[(\alpha \to \beta) \to \gamma) \to \{[(\beta \to \alpha) \to \gamma] \to \gamma\}$$

which is not anymore intuitionistically valid, but it is still forced in any linear Kripke model.)

The Intuitionistic Completeness Theorem shows that if $\Gamma \vdash_{\mathcal{N}} \alpha$ fails, then there is an *uncountable* world in which all formulas in Γ are forced, but α is not. The next result improves on this in two ways: the world is now both *countable* and a *tree*.

Proposition 2.2.9 Countable Model Property. *For any Γ there is a countable world $\mathcal{A}_\Gamma$ in which all formulas of Γ are forced and such that, for every α, if $\Gamma \vdash_{\mathcal{N}} \alpha$ fails, then α is not forced in $\mathcal{A}_\Gamma$.*

Proof. The proof of the Intuitionistic Completeness Theorem 2.2.6 builds a model whose states of knowledge consist of the sets of formulas closed under $\vdash_{\mathcal{N}}$ and extending Γ. The variant introduced by the present proof consists in considering only the *finitely generated* extensions of Γ closed under $\vdash_{\mathcal{N}}$, thus drastically cutting down their number (to only countably many).

Since we only look for finitely generated extensions of Γ, we can easily describe them uniformly as follows. Let $\{\alpha_n\}_{n\in\omega}$ be an enumeration of all implicational formulas, and σ, τ, etc. be sequences of 0's and 1's. Then the finite extensions of Γ can be defined as follows:

$$\begin{aligned} \Gamma_\emptyset &= cl_{\mathcal{N}}(\Gamma) \\ \Gamma_{\sigma * \langle 0 \rangle} &= \Gamma_\sigma \\ \Gamma_{\sigma * \langle 1 \rangle} &= cl_{\mathcal{N}}(\Gamma_\sigma \cup \{\alpha_n\}), \end{aligned}$$

where

$$cl_{\mathcal{N}}(\Theta) = \text{the closure of } \Theta \text{ under } \vdash_{\mathcal{N}}.$$

The idea is simply that, by being a 1 or a 0, the n-th digit of σ tells whether α_n is or is not a member of Γ_σ.

The world $\mathcal{A}$ is defined as follows:

$$\mathcal{A} = \langle \{\Gamma_\sigma\}_{\sigma\in\mathcal{S}}, \subseteq, \{\mathcal{A}_{\Gamma_\sigma}\}_{\sigma\in\mathcal{S}} \rangle,$$

where:

1. $\mathcal{S}$ is the set of all sequences of 0's and 1's

2. $\subseteq$ is the usual set-theoretical inclusion relation

3. $\mathcal{A}_{\Gamma_\sigma}$ is the set of propositional letters in Γ_σ, i.e.

$$\mathcal{A}_{\Gamma_\sigma} = \{p : p \in \Gamma_\sigma\}.$$

As in 2.2.6, we can show that *forcing is reduced to membership*, thanks to closure under $\vdash_{\mathcal{N}}$. For more details, see the proof of the next result. □

We have obtained a definite improvement over the uncountable model used in 2.2.6, but there is more to worry about than simple *cardinality*. In particular, the fact that the construction of the model is not very *effective*. Indeed, at every node we may perform the infinitary operation of taking the closure under $\mathcal{N}$ of a set of formulas (which requires adding, in general, infinitely many consequences). The proof of the next result eliminates this defect.

Proposition 2.2.10 Constructive Model Property. *For any Γ there is a constructively presented world $\mathcal{A}_\Gamma$ in which all formulas of Γ are forced and such that, for every α, if $\Gamma \vdash_{\mathcal{N}} \alpha$ fails, then α is not forced in $\mathcal{A}_\Gamma$.*

Proof. The problem discussed above concerns the use of actual infinity in the construction: the closure operation is infinitary, and it has to be completed before we can go on to the next level. The solution relies on an analysis of the role of closure under $\vdash_{\mathcal{N}}$ in the proof of 2.2.6. What we notice is that such a closure was needed only to be able to replace deducibility from a set of formulas by membership in it. Without closure, we will have to use deducibility directly, but this is not a problem. In other words, instead of proving that forcing is reduced to membership, i.e.

$$\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \alpha \in \Gamma_\sigma,$$

we will just prove that *forcing is reduced to deducibility*, i.e.

$$\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \Gamma_\sigma \vdash_{\mathcal{N}} \alpha.$$

The construction is thus modified as follows:

$$\begin{array}{rcl} \Gamma_\emptyset & = & \Gamma \\ \Gamma_{\sigma*\langle 0\rangle} & = & \Gamma_\sigma \\ \Gamma_{\sigma*\langle n\rangle} & = & \Gamma_\sigma \cup \{\alpha_n\}, \end{array}$$

where $\{\alpha_n\}_{n\in\omega}$ is an enumeration of all implicational formulas.

The world $\mathcal{A}$ is defined as follows:

$$\mathcal{A} = \langle \{\Gamma_\sigma\}_{\sigma\in\mathcal{S}}, \subseteq, \{\mathcal{A}_{\Gamma_\sigma}\}_{\sigma\in\mathcal{S}} \rangle,$$

where:

1. $\mathcal{S}$ is the set of all sequences of 0's and 1's
2. $\subseteq$ is the usual set-theoretical inclusion relation
3. $\mathcal{A}_{\Gamma_\sigma}$ is the set of propositional letters deducible from Γ_σ, i.e.

$$\mathcal{A}_{\Gamma_\sigma} = \{p : \Gamma_\sigma \vdash_{\mathcal{N}} p\}.$$

Notice the change in the definition of $\mathcal{A}_{\Gamma_\sigma}$, which is needed to handle the atomic case in the proof of the reduction of forcing to deducibility.

We can now prove that, for any formula α and any string σ,

$$\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \Gamma_\sigma \vdash_{\mathcal{N}} \alpha.$$

For $\alpha = p$, this holds by the definition of forcing and of $\mathcal{A}_{\Gamma_\sigma}$. For $\alpha = \gamma \to \delta$ it is proved inductively as follows, using →-introduction in one direction and →-elimination in the other:

- Suppose $\Gamma_\sigma \Vdash_{\mathcal{A}} \gamma \to \delta$, i.e.

  $$(\forall \Gamma_\tau \supseteq \Gamma_\sigma)(\Gamma_\tau \Vdash_{\mathcal{A}} \gamma \;\Rightarrow\; \Gamma_\tau \Vdash_{\mathcal{A}} \delta).$$

  By the induction hypothesis,

  $$(\forall \Gamma_\tau \supseteq \Gamma_\sigma)(\Gamma_\tau \vdash_{\mathcal{N}} \gamma \;\Rightarrow\; \Gamma_\tau \vdash_{\mathcal{N}} \delta).$$

  We want $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma \to \delta$. Suppose $\Gamma_\sigma \nvdash_{\mathcal{N}} \gamma \to \delta$. By *→-introduction*, $\Gamma_\sigma, \gamma \nvdash_{\mathcal{N}} \delta$. Let n be such that $\gamma = \alpha_n$, and τ be an extension of σ such that $\Gamma_\tau = \Gamma_\sigma \cup \{\gamma\}$: then $\Gamma_\tau \supseteq \Gamma_\sigma$ and thus, by the induction hypothesis, if $\Gamma_\tau \vdash_{\mathcal{N}} \gamma$, then $\Gamma_\tau \vdash_{\mathcal{N}} \delta$. But $\Gamma_\tau \vdash_{\mathcal{N}} \gamma$ and $\Gamma_\tau \nvdash_{\mathcal{N}} \delta$ by definition of Γ_τ, contradiction.

- Suppose $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma \to \delta$. We want $\Gamma_\sigma \Vdash_{\mathcal{A}} \gamma \to \delta$, i.e.

  $$(\forall \Gamma_\tau \supseteq \Gamma_\sigma)(\Gamma_\tau \Vdash_{\mathcal{A}} \gamma \;\Rightarrow\; \Gamma_\tau \Vdash_{\mathcal{A}} \delta)$$

  or, by the induction hypothesis,

  $$(\forall \Gamma_\tau \supseteq \Gamma_\sigma)(\Gamma_\tau \vdash_{\mathcal{N}} \gamma \;\Rightarrow\; \Gamma_\tau \vdash_{\mathcal{N}} \delta).$$

  Let $\Gamma_\tau \supseteq \Gamma_\sigma$, in particular $\Gamma_\tau \vdash_{\mathcal{N}} \gamma \to \delta$. If $\Gamma_\tau \vdash_{\mathcal{N}} \gamma$, then $\Gamma_\tau \vdash_{\mathcal{N}} \delta$ by *→-elimination*.

Suppose now that $\Gamma \nvdash_{\mathcal{N}} \alpha$. Then, since forcing coincides with deducibility, $\Gamma_\emptyset = \Gamma$ forces every formula in Γ but not α, and thus $\Gamma \not\models_i \alpha$. □

The next result provides a further improvement on the previous ones, by showing that *finite worlds are enough*.

Proposition 2.2.11 Finite Model Property (Smorinsky [1973]) *For any Γ and α there is a finite world $\mathcal{A}_{\Gamma,\alpha}$ such that if $\Gamma \vdash_{\mathcal{N}} \alpha$ fails, then all formulas of Γ are forced, but α is not.*

Proof. Since forcing for a given formula only involves forcing on its subformulas, states that force exactly the same subformulas of (formulas in) Γ and α are indistinguishable from the point of view of Γ and α, and can be collapsed. But there are only finitely many (sets of) subformulas of Γ and α, and thus only finitely many collapsed states.

Formally, given a world

$$\mathcal{A} = \langle P_{\mathcal{A}}, \sqsubseteq_{\mathcal{A}}, \{\mathcal{A}_\sigma\}_{\sigma \in P_{\mathcal{A}}} \rangle$$

we consider the set $S_{\Gamma,\alpha}$ of all subformulas of Γ and α, and define the collapse

$$\mathcal{B} = \langle P_{\mathcal{B}}, \sqsubseteq_{\mathcal{B}}, \{\mathcal{B}_{[\sigma]}\}_{[\sigma] \in P_{\mathcal{B}}} \rangle$$

of $\mathcal{A}$ w.r.t. $S_{\Gamma,\alpha}$ as follows:

1. $P_{\mathcal{B}}$ is the set of equivalence classes w.r.t. to the equivalence relation that identifies states forcing the same subformulas, i.e. the set of all

$$[\sigma] = \{\tau : (\forall \beta \in S_{\Gamma,\alpha})(\sigma \Vdash_{\mathcal{A}} \beta \Leftrightarrow \tau \Vdash_{\mathcal{A}} \beta)\},$$

   for $\sigma \in P_{\mathcal{A}}$. In particular, each equivalence class corresponds to a uniquely determined subset of $S_{\Gamma,\alpha}$.

2. Equivalence classes are ordered as the associated subsets of $S_{\Gamma,\alpha}$:

$$\begin{aligned} [\sigma] \sqsubseteq_{\mathcal{B}} [\tau] \quad &\Leftrightarrow \quad (\forall \beta \in S_{\Gamma,\alpha})(\sigma \Vdash_{\mathcal{A}} \beta \Rightarrow \tau \Vdash_{\mathcal{A}} \beta) \\ &\Leftrightarrow \quad \tau \text{ forces all subformulas forced by } \sigma. \end{aligned}$$

3. Knowledge is preserved in the collapse:

$$\mathcal{B}_{[\sigma]} = \{p : \sigma \Vdash_{\mathcal{A}} p\} = \mathcal{A}_\sigma.$$

Since forcing is determined by state knowledge, it is not surprising that it is preserved by the collapse: for all $\beta \in S_{\Gamma,\alpha}$,

$$\sigma \Vdash_{\mathcal{A}} \beta \Leftrightarrow [\sigma] \Vdash_{\mathcal{B}} \beta.$$

This is proved by induction on β. The atomic case is trivial, since

$$\sigma \Vdash_{\mathcal{A}} p \Leftrightarrow p \in \mathcal{A}_\sigma \Leftrightarrow p \in \mathcal{B}_{[\sigma]} \Leftrightarrow [\sigma] \Vdash_{\mathcal{B}} p.$$

For $\beta = \gamma \rightarrow \delta$ we have:

- $\sigma \Vdash_{\mathcal{A}} \gamma \to \delta \;\Rightarrow\; [\sigma] \Vdash_{\mathcal{B}} \gamma \to \delta$
  Let $\sigma \Vdash_{\mathcal{A}} \gamma \to \delta$, i.e.
  $$(\forall \tau \sqsupseteq_{\mathcal{A}} \sigma)(\tau \Vdash_{\mathcal{A}} \gamma \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \delta).$$
  We want $[\sigma] \Vdash_{\mathcal{B}} \gamma \to \delta$, i.e.
  $$(\forall [\tau] \sqsupseteq_{\mathcal{B}} [\sigma])([\tau] \Vdash_{\mathcal{B}} \gamma \;\Rightarrow\; [\tau] \Vdash_{\mathcal{B}} \delta).$$
  From $[\sigma] \sqsubseteq_{\mathcal{B}} [\tau]$ and $\sigma \Vdash_{\mathcal{A}} \gamma \to \delta$ we have $\tau \Vdash_{\mathcal{A}} \gamma \to \delta$ (by definition of $\sqsubseteq_{\mathcal{B}}$), in particular
  $$\tau \Vdash_{\mathcal{A}} \gamma \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \delta.$$
  By the induction hypothesis,
  $$[\tau] \Vdash_{\mathcal{B}} \gamma \;\Rightarrow\; [\tau] \Vdash_{\mathcal{B}} \delta.$$
  Since $[\tau]$ is an arbitrary extension of $[\sigma]$, $[\sigma] \Vdash_{\mathcal{B}} \gamma \to \delta$.

- $[\sigma] \Vdash_{\mathcal{B}} \gamma \to \delta \;\Rightarrow\; \sigma \Vdash_{\mathcal{A}} \gamma \to \delta$
  Let $[\sigma] \Vdash_{\mathcal{B}} \gamma \to \delta$, i.e.
  $$(\forall [\tau] \sqsupseteq_{\mathcal{B}} [\sigma])([\tau] \Vdash_{\mathcal{B}} \gamma \;\Rightarrow\; [\tau] \Vdash_{\mathcal{B}} \delta).$$
  We want $\sigma \Vdash_{\mathcal{A}} \gamma \to \delta$, i.e.
  $$(\forall \tau \sqsupseteq_{\mathcal{A}} \sigma)(\tau \Vdash_{\mathcal{A}} \gamma \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \delta).$$
  From $\sigma \sqsubseteq_{\mathcal{A}} \tau$ we have, by monotonicity of forcing, $[\sigma] \sqsubseteq_{\mathcal{B}} [\tau]$ (since everything forced by σ is also forced by τ). Thus
  $$[\tau] \Vdash_{\mathcal{B}} \gamma \;\Rightarrow\; [\tau] \Vdash_{\mathcal{B}} \delta$$
  and, by the induction hypothesis,
  $$\tau \Vdash_{\mathcal{A}} \gamma \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \delta.$$
  Since τ is an arbitrary extension of σ, $\sigma \Vdash_{\mathcal{A}} \gamma \to \delta$.

Suppose now that $\Gamma \vdash_{\mathcal{N}} \alpha$ fails. By the Intuitionistic Completeness Theorem, there is a world $\mathcal{A}$ in which all formulas of Γ are forced, but α is not. Then the same happens in the collapse $\mathcal{B}$, which is a finite world. □

The last result is the best possible, since 18.7.2 shows that it is not possible to improve it by making $\mathcal{A}_{\Gamma,\alpha}$ dependent only on Γ, and independent of α. In other words, there is in general no *single* finite world $\mathcal{A}_\Gamma$ in which every formula in Γ is forced, and every formula α such that $\Gamma \vdash_{\mathcal{N}} \alpha$ fails is not.

Actually, 18.7.2 will prove that there is not even a single finite world in which every formula α such that $\vdash_{\mathcal{N}} \alpha$ fails is not forced. However, by combining 2.2.8.a, 2.2.11 and 2.2.9 we get a *countable family* $\{T_n\}_{n\in\omega}$ *of finite trees*, as well as a single *countable tree* T_ω, that do the job. It is possible to just let T_1 be a tree with a single node, T_{n+1} be the tree obtained by adding a single node on top of $n+1$ copies of T_n, and T_ω be the inverse limits of the T_n's, but we do not spell this out because in the following we will get more conspicuous algebraic results of the same kind.

2.3 Tableaux Semantics $\star$

Since the notion of forcing is defined classically, we can use classical methods to investigate it. In particular, we can rephrase everything in terms of classical tableaux (see 19.1.5).

The idea is the following. First, note that we can restrict attention to worlds $\mathcal{A}$ with a least element $0_{\mathcal{A}}$, by adding $0_{\mathcal{A}}$ if it does not exist already, and defining $\mathcal{A}_{0_{\mathcal{A}}}$ as the intersection of the $\mathcal{A}_\sigma$, for all minimal elements $\sigma \in P_{\mathcal{A}}$. Then $0_{\mathcal{A}}$ forces exactly the formulas that are forced by every minimal element σ.

The advantage of having a least element is that, by monotonicity of forcing, if a formula is not forced in $\mathcal{A}$, then it is not forced already by $0_{\mathcal{A}}$. Given α, to determine whether α is intuitionistically valid we can use the definition of forcing to look for worlds $\mathcal{A}$ in which $0_{\mathcal{A}} \Vdash_{\mathcal{A}} \alpha$ fails. If the search is systematic, then either we discover that it is impossible to falsify the condition $0_{\mathcal{A}} \Vdash_{\mathcal{A}} \alpha$ (and then α is intuitionistically valid), or we find a world in which α is not forced.

The rules for a systematic search derive from an analysis of the forcing definition:

$$\sigma \Vdash_{\mathcal{A}} \alpha \to \beta \;\Leftrightarrow\; (\forall \tau \sqsupseteq_{\mathcal{A}} \sigma)(\tau \Vdash_{\mathcal{A}} \alpha \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \beta),$$

where the implication on the right is classical. In other words:

1. If $\sigma \Vdash_{\mathcal{A}} \alpha \to \beta$ is true, then

   $$\tau \Vdash_{\mathcal{A}} \alpha \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \beta$$

   holds for *any* extension τ of σ. By the classical meaning of implication, either $\tau \Vdash_{\mathcal{A}} \alpha$ fails or $\tau \Vdash_{\mathcal{A}} \beta$ holds, for *any* extension τ of σ.

2. If $\sigma \Vdash_{\mathcal{A}} \alpha \to \beta$ is false, then

   $$\tau \Vdash_{\mathcal{A}} \alpha \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \beta$$

   fails for *some* extension τ of σ. By the classical meaning of implication, $\tau \Vdash_{\mathcal{A}} \alpha$ must hold and $\tau \Vdash_{\mathcal{A}} \beta$ must fail, for *some* extension τ of σ.

Intuitionistic tableaux

We now turn the previous remarks into formal rules.

Definition 2.3.1 *An* **intuitionistic tableau** *is a tree with nodes consisting of signed forcing assertions of the form $T\sigma \Vdash \alpha$ or $F\sigma \Vdash \alpha$, and consistent with the following formation rules:*

1. *If a node $T\sigma \Vdash \alpha \to \beta$ is on the tree, then we can split any branch going through it by adding $F\tau \Vdash \alpha$ in one direction and $T\tau \Vdash \beta$ in the other, where τ is any extension of σ that has already been introduced. Graphically,*

$$\begin{array}{c} T\sigma \Vdash \alpha \to \beta \\ \hline\hline F\tau \Vdash \alpha \qquad T\tau \Vdash \beta, \end{array}$$

where the double horizontal line shows that the bottom nodes do not have to immediately follow the top one.

2. *If a node $F\sigma \Vdash \alpha \to \beta$ is on the tree, then we can extend any branch going through it by adding $T\tau \Vdash \alpha$ and $F\tau \Vdash \beta$, where τ is a new extension of σ (incomparable with all other extensions of σ already introduced on the same branch). Graphically,*

$$\begin{array}{c} F\sigma \Vdash \alpha \to \beta \\ \hline\hline T\tau \Vdash \alpha \\ \hline F\tau \Vdash \beta. \end{array}$$

Notice the asymmetric treatment: in the first case we split branches and consider extensions of σ that have already been introduced; in the second case we linearly extend branches and introduce new extensions of σ. The reason for the different treatment of the extension of σ is quite obvious: in the first case we consider only the states already introduced, leaving open the possibility of considering new ones in the future, if needed; in the second case we only know that some extension of σ has the required property, and we introduce a new one because there is no reason to believe that any of the ones already introduced is the appropriate one.

Intuitionistic tableaux provide an alternative approach to provability, with a strong semantical flavor.

Definition 2.3.2 *A formula α is* **provable by intuitionistic tableaux** *from Γ (written $\boldsymbol{\Gamma \vdash_{\mathcal{T}} \alpha}$) if there is an intuitionistic tableau starting from the nodes $T0 \Vdash \gamma$, for all $\gamma \in \Gamma$, and $F0 \Vdash \alpha$ such that all its branch are contradictory, in the sense that on every branch there is a pair of nodes of the form $T\sigma \Vdash \beta$ and $F\tau \Vdash \beta$, with σ and τ compatible (the same β for any given branch, although possibly different β's for different branches).*

The use of different compatible states σ and τ can be avoided, if we add a monotonicity rule as follows:

$$T\sigma \Vdash \alpha \wedge \sigma \sqsubseteq \tau \implies T\tau \Vdash \alpha.$$

The next definition captures the idea of systematic search.

Definition 2.3.3 *A* **complete systematic tableau** *is a tableau in which the rules have been used exhaustively, in the sense that:*

1. *For any node $T\sigma \Vdash \alpha \to \beta$ on the tree, any branch going through it and any extension τ of σ used in the tree, there is a node on the branch that splits into two branches going through the nodes $F\tau \Vdash \alpha$ and $T\tau \Vdash \beta$.*

2. *For any node $F\sigma \Vdash \alpha \to \beta$ on the tree and any branch going through it, there is a node on the branch followed by two nodes $T\tau \Vdash \alpha$ and $F\tau \Vdash \beta$, where τ is a new extension of σ (incomparable with all other extensions of σ that previously occur in the given branch).*

In an actual construction of a complete systematic tableau, nodes of the type $F\sigma \Vdash \alpha \to \beta$ need be considered only once, by extending every branch going through them in the appropriate way (in particular, by introducing a new extension of σ for every such branch). Nodes of the type $T\sigma \Vdash \alpha \to \beta$ cannot instead be considered once and for all, since new extensions of σ may be introduced in the following: we thus need to periodically go back to such nodes, and consider all newly introduced extensions of σ. In particular, *a complete systematic tableau may be infinite*.

The procedure can obviously be carried through in an orderly fashion, thus producing a complete systematic tableau, starting from any given signed forcing assertion.

Obviously, whenever in the construction of a complete systematic intuitionistic tableau we hit a contradiction along a branch, we can seal that branch off and stop developing it, since every extension of it will remain contradictory.

Also, as partial orderings we only really need to consider suborderings of a sufficiently universal one, e.g. the set of all finite strings of natural numbers ordered lexicographically.

Soundness and Completeness

Since the notion of provability by tableaux was modelled on semantical notions, the connections with semantics are particular transparent and easy to prove. We first give two examples, to illustrate both a success and a failure, and then establish the connection in the form of a Soundedness and Completeness Theorem.

The first example shows that Axiom 3 of 1.2.3 is provable by intuitionistic tableaux.

$$\begin{array}{c}
F0 \Vdash [p \to (q \to r)] \to [(p \to q) \to (p \to r)] \\ \hline
T00 \Vdash p \to (q \to r) \\ \hline
F00 \Vdash (p \to q) \to (p \to r) \\ \hline
T000 \Vdash p \to q \\ \hline
F000 \Vdash p \to r \\ \hline
T0000 \Vdash p \\ \hline
F0000 \Vdash r \\ \hline
\begin{array}{cc}
F0000 \Vdash p & \begin{array}{c}
T0000 \Vdash q \\ \hline
\begin{array}{cc}
F0000 \Vdash p & \begin{array}{c}
T0000 \Vdash q \to r \\ \hline
\begin{array}{cc}
F0000 \Vdash q & T0000 \Vdash r.
\end{array}
\end{array}
\end{array}
\end{array}
\end{array}
\end{array}$$

The linear top part comes from the analysis of the three false forcing conditions, as they arise: each of them introduces a new extension. Then we analyze $T000 \Vdash p \to q$ and $T00 \Vdash p \to (q \to r)$, both times using 0000 as an extension (of 000 and 00, respectively), since we can choose any extension. This produces contradictions on every branch.

The second example shows that a complete systematic tableau for Peirce's Law does not close.

$$\begin{array}{c}
F0 \Vdash [(p \to q) \to p] \to p \\ \hline
T00 \Vdash (p \to q) \to p \\ \hline
F00 \Vdash p \\ \hline
\begin{array}{cc}
\begin{array}{c}
F00 \Vdash p \to q \\ \hline
T000 \Vdash p \\ \hline
F000 \Vdash q \\ \hline
\begin{array}{cc}
\begin{array}{c}
F000 \Vdash p \to q \\ \hline
T0000 \Vdash p \\ \hline
F0000 \Vdash q \\ \hline
\cdots
\end{array} & T000 \Vdash q
\end{array}
\end{array} & T00 \Vdash p
\end{array}
\end{array}$$

Note that all branches branching right are contradictory, but the tree continues to grow in a periodical way, since nodes $F\sigma \Vdash p \to q$ continue to introduce new extension of 00, and this forces us to go back to the node $T00 \Vdash (p \to q) \to p$.

The different behavior on the two examples is connected to the fact that the first formula is intuitionistically valid, while the second is not. The connection is the content of the next results.

Proposition 2.3.4 Soundness Theorem for Tableaux (Nerode [1990]) *For any* Γ *and* α,

$$\Gamma \vdash_{\mathcal{T}} \alpha \;\Rightarrow\; \Gamma \models_i \alpha.$$

Proof. We prove the contrapositive, i.e. if there is a world $\mathcal{A}$ in which all formulas in Γ are forced but α is not, then any intuitionistic tableau starting from $T0 \Vdash \gamma$, for all $\gamma \in \Gamma$, and $F0 \Vdash \alpha$ has a noncontradictory branch.

This is achieved by showing that if a world $\mathcal{A}$ agrees with the initial verteces of an intuitionistic tableau, then it must agree with a branch of it, in the sense that there is a way of interpreting states σ named on the branch by states $\sigma_{\mathcal{A}} \in P_{\mathcal{A}}$ in such a way that

$$\begin{array}{lcl} T\sigma \Vdash \alpha \text{ is on the branch} & \Rightarrow & \sigma_{\mathcal{A}} \Vdash_{\mathcal{A}} \alpha \\ F\sigma \Vdash \alpha \text{ is on the branch} & \Rightarrow & \sigma_{\mathcal{A}} \not\Vdash_{\mathcal{A}} \alpha. \end{array}$$

This is easily proved by induction on the construction of the tree: the initial verteces agree with $\mathcal{A}$ by choice of $\mathcal{A}$ (by adding a smallest element $0_{\mathcal{A}}$ to $P_{\mathcal{A}}$ if it does not exist already), and the rules for the construction of tableaux were chosen to mirror the definition of forcing. The only case that requires a small argument is $F\sigma \Vdash \alpha \to \beta$: if $T\tau \Vdash \alpha$ and $F\tau \Vdash \beta$ are on the tree, it is enough to let $\tau_{\mathcal{A}}$ be any extension of $\sigma_{\mathcal{A}}$ such that $\tau_{\mathcal{A}} \Vdash_{\mathcal{A}} \alpha$ and $\tau_{\mathcal{A}} \not\Vdash_{\mathcal{A}} \beta$, and such a state exists because, by the induction hypothesis, $\sigma_{\mathcal{A}} \not\Vdash_{\mathcal{A}} \alpha \to \beta$.

Since a world cannot at the same time force and not force a formula, a branch agreeing with a world cannot be contradictory. We have thus proved that no intuitionistic tableau starting from $T0 \Vdash \gamma$ (for all $\gamma \in \Gamma$) and $F0 \Vdash \alpha$ has a noncontradictory branch. $\Box$

Proposition 2.3.5 Completeness Theorem for Tableaux (Nerode [1990]) *For any Γ and α,*

$$\Gamma \models_i \alpha \;\Rightarrow\; \Gamma \vdash_{\mathcal{T}} \alpha.$$

Proof. We prove the contrapositive: if a branch of a complete systematic tableau starting from $T0 \Vdash \gamma$ (for all $\gamma \in \Gamma$) and $F0 \Vdash \alpha$ is noncontradictory, then there is a world $\mathcal{A}$ in which all formulas in Γ are forced but α is not.

This is achieved by building a world $\mathcal{A}$ that agrees with the branch, in the sense of the previous proposition. We only have to use the information on the branch itself, and turn it into a world. We thus let:

1. $P_{\mathcal{A}}$ be the set of all states σ appearing on some node of the branch, in particular $\sigma_{\mathcal{A}} = \sigma$

2. $\sqsubseteq_{\mathcal{A}}$ be the lexicographical order on $P_{\mathcal{A}}$

3. $\mathcal{A}_\sigma$ be the set of all p such that $T\tau \Vdash p$ is on the branch, for some $\tau \sqsubseteq_{\mathcal{A}} \sigma$.

The consideration of all $\tau \sqsubseteq_{\mathcal{A}} \sigma$, instead of σ alone, ensures the required monotonicity of $\mathcal{A}_\sigma$.

We prove, by induction of formulas, that $\mathcal{A}$ agrees with the given branch. If $T\sigma \Vdash p$ is on the branch then $p \in \mathcal{A}_\sigma$ by definition, and thus $\sigma \Vdash_{\mathcal{A}} p$. If $F\sigma \Vdash p$ is on the branch, then $T\sigma \Vdash p$ is not, because the branch is noncontradictory: then $p \notin \mathcal{A}_\sigma$, and $\sigma \not\Vdash_{\mathcal{A}} p$. The remaining cases are true by induction, since the rules for the construction of tableaux were chosen to mirror the definition of forcing.

We have thus proved that there is a world $\mathcal{A}$ that forces all formulas in Γ and does not force α, since the verteces $T0 \Vdash \gamma$ (for any $\gamma \in \Gamma$) and $F0 \Vdash \alpha$ are on every branch, in particular on the noncontradictory one we considered. $\Box$

æ

Chapter 3

Complexity

Historically, the first presentations of the implicational calculus (actually, of the full classical propositional calculus) were Hilbert Systems. They had the obvious disadvantage of being based on the single rule of Modus Ponens, which does not allow for a backward search of a proof of a theorem, because the premises contain a formula unrelated to the conclusion. Proofs in the Hilbert Systems were usually difficult to find, and cumbersome to write down.

Natural Deduction is a substantial improvement of Hilbert Systems, in which axioms are reduced to trivial facts and replaced by rules. The disadvantage of having Modus Ponens among the rules however remains, and again this does not allow for a straightforward backward search of a proof. Normal proofs however have the subformula property, and thus the normalization theorem shows that we can restrict the search for possible proofs to normal ones, thus providing a decision procedure.

The Sequent System has a built-in subformula property, which is destroyed by the introduction of the Cut Rule for the purpose of proving equivalence with Natural Deduction, and restored by the Cut Elimination Theorem. Working with sequents provides the simplest decision procedure.

In this chapter we will first spell out various decision procedures for the implicational calculus, and then measure their complexity.

3.1 Decision Procedures

The decision procedure for $\mathcal{N}$ is a consequence of the Subformula Property for normal proofs (1.1.2.2). But some care is needed since, despite the Subformula Property, *a formula can have infinitely many different normal proofs*. For example, in the following normal proof of the formula $(p \to p) \to (p \to p)$:

$$\dfrac{\dfrac{\dfrac{\dfrac{[p]^{(1)} \quad [p \to p]^{(2)}}{p} \qquad [p \to p]^{(2)}}{p}}{p^{(1)} \to p}}{(p \to p)^{(2)} \to (p \to p),}$$

we can repeat the initial part (here repeated twice) any finite number of times.

Proposition 3.1.1 *The relation $\vdash_{\mathcal{N}}$ is decidable.*

Proof. By 1.1.3, every proof can be transformed into one in normal form. Given Γ and α, it thus suffices to look for a proof of $\Gamma \vdash_{\mathcal{N}} \alpha$ in normal form. By 1.1.2.2, only subformulas of Γ and α can occur in such a proof: since there are only finitely many such subformulas, there are only finitely many possible *irredundant* proofs of $\Gamma \vdash_{\mathcal{N}} \alpha$, i.e. proofs in which the same formula never appears twice on the same branch.

Notice that if a proof is not irredundant, it can be transformed into an irredundant one (since what happens on the tree between two repetitions of the same formula is inessential). It is thus enough to look for an irredundant proof, by systematically checking all possible irredundant trees built from subformulas of Γ and α according to the rules of $\mathcal{N}$. If a proof is found, then $\Gamma \vdash_{\mathcal{N}} \alpha$ is provable. Otherwise, it is not. □

The decision procedure for $\mathcal{S}$ is again a consequence of the Subformula Property. Again some care is needed because, as for normal proofs, *a sequent can have infinitely many different cut-free proofs*. For example, in the following cut-free proof of the sequent $\vdash_{\mathcal{S}} (p \to p) \to (p \to p)$:

$$\dfrac{\dfrac{\dfrac{\dfrac{p \vdash_{\mathcal{S}} p \quad p \vdash_{\mathcal{S}} p}{p, p \to p \vdash_{\mathcal{S}} p} \qquad p \vdash_{\mathcal{S}} p}{p, p \to p \vdash_{\mathcal{S}} p}}{p \to p \vdash_{\mathcal{S}} p \to p}}{\vdash_{\mathcal{S}} (p \to p) \to (p \to p),}$$

we can repeat the initial part (here repeated twice) any finite number of times.

Proposition 3.1.2 *The relation $\vdash_{\mathcal{S}}$ is decidable.*

Proof. Given a sequent $\Gamma \vdash_{\mathcal{S}} \alpha$, by the Subformula Property only subformulas of Γ and α can occur in any of its proofs: since there are only finitely many such subformulas, there are only finitely many possible *irredundant* proofs of $\Gamma \vdash_{\mathcal{S}} \alpha$, i.e. proofs in which the same sequent never appears twice on the same branch.

Notice that if a proof is not irredundant, it can be transformed into an irredundant one (since what happens on the tree between two repetitions of the same sequent is inessential). It is thus enough to look for an irredundant proof, by systematically checking all possible irredundant trees built from subformulas of Γ and α according to the rules of $\mathcal{S}$. If a proof is found, then $\Gamma \vdash_{\mathcal{S}} \alpha$ is provable. Otherwise, it is not. □

The decision procedure for $\models_i$ is a consequence of the Finite Model Property 2.2.11.

Proposition 3.1.3 *The relation $\models_i$ is decidable.*

Proof. The Finite Model Property shows that, for any given Γ and α, either $\Gamma \vdash_{\mathcal{N}} \alpha$ or there is a finite world in which all formulas of Γ are forced, but α is not. We can thus simultaneously generate all proofs of $\mathcal{N}$ and all finite worlds, until one of the two cases happens. □

Among the decision procedures proposed above, the most efficient is the one based on sequents: we only have to start from the sequent $\Gamma \vdash_{\mathcal{S}} \alpha$ and work our way up, by systematically pursuing every possible $\rightarrow$-introduction, both on the left and on the right. Notice that, unlike in the classical case 19.3.1, *not every possible analysis of $\rightarrow$-introduction produces a proof*, even when a proof does exist. For example,

$$p, p \rightarrow q, s \rightarrow t \vdash_{\mathcal{S}} q$$

is provable, but the following is not a proof:

$$\dfrac{\dfrac{p \vdash_{\mathcal{S}} p \qquad q \vdash_{\mathcal{S}} s}{p, p \rightarrow q \vdash_{\mathcal{S}} s} \qquad t \vdash_{\mathcal{S}} q}{p, p \rightarrow q, s \rightarrow t \vdash_{\mathcal{S}} q.}$$

Thus we really have to consider every possible analysis, before claiming that a formula is not provable.

3.2 Complexity of Decision Procedures

Computational Complexity

Theorem 3.2.1 (Statman [1976]) *The complexity of $\models_i$ is PSPACE- complete.*

Proof.
□

Complexity of Normalization

Proposition 3.2.2 *A proof of height bounded by h and with maxima on formulas of degree at most n can be replaced by an equivalent proof of height bounded by 2^h and maxima on formulas of degree at most $n-1$.*

Proof.

□

Corollary 3.2.3 Normalization Theorem. *A proof of height bounded by h and with maxima on formulas of degree at most n can be replaced by an equivalent normal proof of height bounded by $2^{2^{\cdots 2^{h}}}$, with n iterations of the exponential.*

Proof.

□

Complexity of Cut Elimination

Proposition 3.2.4 *A proof of height bounded by h and with cuts on formulas of degree at most n can be replaced by an equivalent proof of height bounded by 2^h and cuts on formulas of degree at most $n-1$.*

Proof.

□

Corollary 3.2.5 Cut Elimination Theorem. *A proof of height bounded by h and with cuts on formulas of degree at most n can be replaced by an equivalent cut-free proof of height bounded by $2^{2^{\cdots 2^{h}}}$, with n iterations of the exponential.*

Proof.

□

The effect of the elimination of cuts is thus an explosion of the size of the proof. Equivalently, the advantage of using cuts is the possibility of giving compact proofs. æ

Chapter 4

Conjunction

In this chapter we introduce the new connective of *conjunction*. On the one hand, the present extension is uncontroversial and mild. On the other hand, however, it allows for a full algebraic treatment of implication via *Heyting $\sqcap$-algebras* (see Chapter 5), as well as for a complete description of models via *cartesian closed categories* (see Chapter 6).

The main parallels among the various notions introduced in the book is succinctely stated in the following table:

Logic (Ch. 4)	Algebra (Ch. 5)	Categories (Ch. 6)
$\alpha \vdash_{\mathcal{N}} \beta$	$a \sqsubseteq b$	$\mathit{Hom}(A, B) \neq \emptyset$
$\alpha \to \beta$	$a \Rightarrow b$	$A \Rightarrow B$
$\alpha \wedge \beta$	$a \sqcap b$	$A \times B$
$\underline{\alpha \wedge \beta \vdash_{\mathcal{N}} \gamma}$	$\underline{(a \sqcap b) \sqsubseteq c}$	$\underline{\mathit{Hom}(A \times B, C) \neq \emptyset}$
$\alpha \vdash_{\mathcal{N}} \beta \to \gamma$	$a \sqsubseteq (b \Rightarrow c)$	$\mathit{Hom}(A, B \Rightarrow C) \neq \emptyset$
T(true)	1 (greatest element)	1 (terminal object).

Additional extensions of the Implicational Calculus with Conjunction are treated in Chapter 17, to which the interested reader can turn immediately after reading the present chapter.

Implicational Calculus with Conjunction

We extend Implicational Calculus as follows:

1. the **language** has an added connective $\wedge$ (*conjunction*)

2. the definition of **formulas** has an added inductive clause, i.e.

    - if α and β are formulas, so is $(\alpha \wedge \beta)$.

To increase readability, some parentheses can be omitted according to the precedence rule: *conjunction over implication*. For example, the formula

$$\alpha \wedge \beta \to \gamma$$

that, as it stands, would be ambiguous, will be read as

$$(\alpha \wedge \beta) \to \gamma$$

and not as

$$\alpha \wedge (\beta \to \gamma).$$

The goal of this chapter is to determine which of the formulas of the Implicational Calculus with Conjunction can be considered 'true', when the connective $\wedge$ is intuitively taken as representing 'conjunction'.

Following the blueprint of Chapters 1–3, we introduce different analyses and study their equivalence and complexity. We continue to use the same symbols $\vdash_{\mathcal{N}}$, $\vdash_{\mathcal{H}}$, $\vdash_{\mathcal{S}}$, $\vdash_{\mathcal{T}}$ and $\models_i$, but they now refer to the extended system with implication and conjunction.

4.1 Syntax

Natural Deduction

The extension of Natural Deduction to conjunction is unproblematic.

Definition 4.1.1 (Gentzen [1935]) *The relation $\vdash_{\mathcal{N}}$ defined in 1.1.1 is extended to conjunction as follows:*

4. **$\wedge$-Introduction**. *If both α and β are deducible from Γ, then so is $\alpha \wedge \beta$:*

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta.}$$

5. **$\wedge$-Elimination**. *If $\alpha \wedge \beta$ is deducible from Γ, then so is any of α and β:*

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta}{\Gamma \vdash_{\mathcal{N}} \alpha} \qquad \textit{and} \qquad \frac{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta}{\Gamma \vdash_{\mathcal{N}} \beta.}$$

An application of $\wedge$-introduction combines two proofs $\mathcal{D}_\alpha$ and $\mathcal{D}_\beta$ into a proof of $\alpha \wedge \beta$. When followed by a $\wedge$-elimination, we obtain a proof of α or β, for example:

$$\dfrac{\dfrac{\begin{array}{c}\Gamma \\ \mathcal{D}_\alpha \\ \alpha\end{array} \qquad \begin{array}{c}\Gamma \\ \mathcal{D}_\beta \\ \beta\end{array}}{\alpha \wedge \beta}}{\alpha.}$$

The occurrence of $\alpha \wedge \beta$ in such a proof is called a **maximum** (relative to $\wedge$).

A more direct way of getting from $\mathcal{D}_\alpha$ and $\mathcal{D}_\beta$ to a proof of α is obviously to forget about $\mathcal{D}_\beta$, and the step

$$\text{from}\qquad \dfrac{\dfrac{\begin{matrix}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{matrix}\qquad\begin{matrix}\Gamma\\ \mathcal{D}_\beta\\ \beta\end{matrix}}{\alpha\wedge\beta}}{\alpha}\qquad\text{to}\qquad \begin{matrix}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{matrix}$$

is called a **maximum elimination**. A symmetric maximum elimination can be obtained by working on β. A proof is in **normal form** if it has no maxima relative to either $\rightarrow$ or $\wedge$.

We can extend the notion of a **descending path** of a normal proof (p. 11) by allowing the path to go through arbitrary $\wedge$-eliminations or $\wedge$-introductions (i.e., we only rule out going through minor premises of $\rightarrow$-eliminations). Then the proof of 1.1.2 actually proves the following.

Proposition 4.1.2 Structure of Normal Proofs (Prawitz [1965]) *For a normal proof of $\mathcal{N}$ the following hold:*

1. **Elimination-Introduction Separation**. *Any descending path consists of two (possibly empty) parts: a first (upper) one going only through eliminations, and a second (lower) one going only through introductions.*

2. **Subformula Property**. *Any formula occurring in the proof is a subformula of either an undischarged assumption or the conclusion.*

Also the Weak Normalization Theorem continues to hold.

Theorem 4.1.3 Weak Normalization (Prawitz [1965]) *Every proof can be transformed into a normal proof, by means of an appropriate sequence of maxima eliminations.*

Proof. As in the case of maxima relative to $\rightarrow$, the elimination of a maximum

$$\dfrac{\dfrac{\begin{matrix}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{matrix}\qquad\begin{matrix}\Gamma\\ \mathcal{D}_\beta\\ \beta\end{matrix}}{\alpha\wedge\beta}}{\alpha}\qquad\text{into}\qquad \begin{matrix}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{matrix}$$

can introduce a new maximum, by turning the final occurrence of α into a maximum (if the last step of the proof $\mathcal{D}_\alpha$ is an introduction of $\rightarrow$ or $\wedge$, and the first one of the rest of the proof below α is an elimination of the same connective). As in 1.1.3, we only need to extend the notion of complexity to take care of the case of conjunction as well, by adding to the definition of **degree** of a formula the following clause:

- the degree of $\alpha \wedge \beta$ is 1 plus the greatest of the degrees of α and β.

The rest of the proof is as in 1.1.3. □

The Normalization Theorem and the last part of the Subformula Property show that no new formulas of the Implicational Calculus can be proved in the extended system with conjunction: if $\wedge$ does not occur in the premises or in the conclusion of a normal proof, then it does not occur at all in the proof. In technical terms, *the system with implication and conjunction is a conservative extension of the system with implication alone*.

Hilbert systems

When dealing with a new connective, we have two choices to extend a Hilbert system: to add either new *axioms*, or new *rules*. Adding rules, however, does not avoid the need of adding new axioms: for example, the three axioms of 1.2.3 were needed to deal with the rule of Modus Ponens. It is thus better to simply add axioms relative to the new connective, and keep Modus Ponens as the only rule. The definition of $\vdash_{\mathcal{H}}$ is thus unchanged, and the Deduction Theorem is still valid as before (since its validity only depends on the presence of Axioms 2 and 3 of 1.2.3, and the fact that Modus Ponens is the only rule).

The axioms for $\wedge$ are chosen in such a way as to make the proof of the equivalence of the systems $\mathcal{N}$ and $\mathcal{H}$ trivial. In particular, they mimic the $\wedge$-introduction and $\wedge$-elimination rules.

Theorem 4.1.4 Equivalence of Hilbert Systems and Natural Deduction (Gentzen [1935]) *If $\mathcal{H}$ is any Hilbert system whose theorems include 1–3 of 1.2.3 and, for any α, β, and γ, the following:*

4. $\alpha \to (\beta \to \alpha \wedge \beta)$

5. $\alpha \wedge \beta \to \alpha$

6. $\alpha \wedge \beta \to \beta$,

then for any Γ and β:

$$\Gamma \vdash_{\mathcal{H}} \beta \ \Leftrightarrow\ \Gamma \vdash_{\mathcal{N}} \beta.$$

Proof. The right to left direction is obtained by induction on the definition of $\vdash_{\mathcal{N}}$. The only cases not dealt with in 1.2.2 are the ones relative to $\wedge$, which we now treat.

Suppose $\Gamma \vdash_{\mathcal{H}} \alpha$ and $\Gamma \vdash_{\mathcal{H}} \beta$. Then we get $\Gamma \vdash_{\mathcal{H}} \alpha \wedge \beta$ by inserting the given proofs of α and β from Γ above their occurrences in the following partial proof (from assumption 4):

$$\frac{\beta \qquad \dfrac{\alpha \qquad \alpha \to (\beta \to \alpha \wedge \beta)}{\beta \to \alpha \wedge \beta}}{\alpha \wedge \beta.}$$

Suppose $\Gamma \vdash_{\mathcal{H}} \alpha \wedge \beta$. Then we get $\Gamma \vdash_{\mathcal{H}} \alpha$ by inserting the given proof of $\alpha \wedge \beta$ from Γ above its occurrence in the following partial proof (from assumption 5):

$$\frac{\alpha \wedge \beta \qquad \alpha \wedge \beta \to \alpha}{\alpha.}$$

Similarly for $\Gamma \vdash_{\mathcal{H}} \beta$, using 6.

For the right to left direction, we only need to show that 4, 5 and 6 are provable in Natural Deduction. 4 is proved by

$$\cfrac{\cfrac{\cfrac{[\alpha]^{(2)} \qquad [\beta]^{(1)}}{\alpha \wedge \beta}}{\beta^{(1)} \to \alpha \wedge \beta}}{\alpha^{(2)} \to (\beta^{(1)} \to \alpha \wedge \beta).}$$

5 is proved by

$$\cfrac{\cfrac{[\alpha \wedge \beta]^{(1)}}{\alpha}}{(\alpha \wedge \beta)^{(1)} \to \alpha.}$$

6 is proved similarly. □

Sequents

As already for Natural Deduction, the extension of the Sequent System to conjunction is unproblematic.

Definition 4.1.5 (Gentzen [1935]) *The relation $\vdash_{\mathcal{S}}$ defined in 1.3.1 is extended to conjunction as follows:*

4. **$\wedge$-Introduction on the right**. *If both α and β are deducible from Γ, then so is $\alpha \wedge \beta$:*

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \wedge \beta.}$$

5. **$\wedge$-Introduction on the left**. *If γ is deducible from Γ, α or Γ, β, then it is deducible from Γ and $\alpha \wedge \beta$:*

$$\frac{\Gamma, \alpha \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \wedge \beta \vdash_{\mathcal{S}} \gamma} \qquad \textit{and} \qquad \frac{\Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \wedge \beta \vdash_{\mathcal{S}} \gamma.}$$

Because of the fact that the Thinning Rule

$$\frac{\Gamma, \vdash_{\mathcal{S}} \gamma}{\Gamma, \Delta \vdash_{\mathcal{S}} \gamma}$$

is a derived rule (see p. 20), we can actually rephrase the two rules of ∧-introduction on the left as a single rule

$$\frac{\Gamma, \alpha, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \wedge \beta \vdash_{\mathcal{S}} \gamma,}$$

which we will use in the following.

As usual, the rules of $\mathcal{S}$ are *backward deterministic* and the **Subformula Property** continues to hold.

The systems $\mathcal{N}$ and $\mathcal{S}$ are obviously equivalent in presence of the Cut Rule. In particular, the translation from $\mathcal{N}$ to $\mathcal{S}$ requires proving the rules of $\mathcal{N}$ as derived rules of $\mathcal{S}$. On the one hand, ∧-introduction of $\mathcal{N}$ corresponds to ∧-introduction on the right. On the other hand, ∧-elimination can be dealt with by ∧-introduction on the left and Cut, as follows:

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \wedge \beta \qquad \dfrac{\Gamma, \alpha, \beta \vdash_{\mathcal{S}} \alpha}{\Gamma, \alpha \wedge \beta \vdash_{\mathcal{S}} \alpha}}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \alpha.}$$

Also the translation from $\mathcal{S}$ to $\mathcal{N}$ poses no problem. On the one hand, ∧-introduction on the right corresponds to ∧-introduction of $\mathcal{N}$. On the other hand, ∧-introduction on the left can be dealt with as follows: given $\Gamma, \alpha, \beta \vdash_{\mathcal{N}} \gamma$, we insert $\alpha \wedge \beta$ above every undischarged occurrence of α and β, which is a permissible step because of ∧-elimination, and this produces a proof from the assumptions Γ and $\alpha \wedge \beta$, i.e. $\Gamma, \alpha \wedge \beta \vdash_{\mathcal{N}} \gamma$.

To get the full equivalence between the two systems, we need to extend 1.3.3.

Theorem 4.1.6 Cut Elimination (Gentzen [1935]) *For any* Γ *and* β*:*

$$\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \beta \;\Rightarrow\; \Gamma \vdash_{\mathcal{S}} \beta.$$

Proof. The Cut Elimination procedure 22, to which we refer in the following, can be extended to take care of the new connective ∧.

A cut was called *inductive* when the formula which is cut has just been introduced on both sides:

$$\frac{\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \wedge \beta} \qquad \dfrac{\Gamma, \alpha, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \wedge \beta \vdash_{\mathcal{S}} \gamma}}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \gamma.}$$

Such a cut can be eliminated as follows, by substituting it with *two* cuts on the formulas α and β (of lower degree):

$$\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \dfrac{\Gamma \vdash_{\mathcal{S}} \beta \quad \Gamma, \alpha, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \vdash_{\mathcal{S}+\mathrm{Cut}} \gamma}}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \gamma.}$$

A cut was called *interlocutory* when the formula which is cut has been introduced at steps preceeding the last ones (on the appropriate sides). In this case we simply move the cut upwards, until it can be eliminated as above. For example, a cut like

$$\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \dfrac{\Gamma, \alpha, \gamma, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \alpha, \gamma \wedge \delta \vdash_{\mathcal{S}} \beta}}{\Gamma, \gamma \wedge \delta \vdash_{\mathcal{S}+\mathrm{Cut}} \beta}$$

can be replaced by one as

$$\dfrac{\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \quad \Gamma, \alpha, \gamma, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \gamma, \delta \vdash_{\mathcal{S}+\mathrm{Cut}} \beta}}{\Gamma, \gamma \wedge \delta \vdash_{\mathcal{S}} \beta.}$$

The remaining cases are similar. □

The Cut Elimination Theorem fills the remaining gap in the proof of the following result.

Corollary 4.1.7 Equivalence of the Sequent System and Natural Deduction (Gentzen [1935]) *For any* Γ *and* β*:*

$$\Gamma \vdash_{\mathcal{S}} \beta \;\Leftrightarrow\; \Gamma \vdash_{\mathcal{N}} \beta.$$

A different proof of Cut Elimination comes from the following extension of 1.3.6.

Proposition 4.1.8 (Prawitz [1965]) *There are canonical translations of cut-free proofs in* $\mathcal{S}$ *to normal proofs in* $\mathcal{N}$*, and conversely.*

Proof. The translation from $\mathcal{S}$ to $\mathcal{N}$ given as half of the proof of the equivalence of the two systems already shows that (cut-free) proofs in $\mathcal{S}$ correspond to normal proofs in $\mathcal{N}$.

For the converse, we only have to supplement 1.3.6. The case of the *introduction rules* of $\mathcal{N}$ can be dealt with by induction and the corresponding introduction rules on the right. For example, if $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ has been obtained by $\wedge$-introduction from normal proofs $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \beta$, then $\Gamma \vdash_{\mathcal{S}} \alpha$ and $\Gamma \vdash_{\mathcal{S}} \beta$ by the induction hypothesis, and thus $\Gamma \vdash_{\mathcal{S}} \alpha \wedge \beta$ by $\wedge$-introduction on the right.

The case of the *elimination rules* of $\mathcal{N}$ is the crucial one, and requires more ingenuity because the natural translation uses the Cut Rule. The general schema has been shown in 1.3.6: if $\Gamma \vdash_{\mathcal{N}} \beta$ has been obtained by an elimination rule, the latter is either a $\rightarrow$-elimination or a $\wedge$-elimination. The two cases are treated similarly, by climbing up in the given proof until an assumption is reached that is *eliminated* in the first step of the given proof below it: this is possible because the given proof is normal, and β has been obtained by an elimination rule. Such an assumption must be of one of the following two forms: $\gamma \rightarrow \delta$ and $\gamma \wedge \delta$. The former case has already been dealt with in 1.3.6, and we now consider the latter. The given proof is then, for example, of the form:

$$\Gamma,\ \begin{array}{c}\dfrac{\gamma \wedge \delta}{\gamma}\\ \mathcal{D}\\ \beta.\end{array}$$

We apply the induction hypothesis to $\mathcal{D}$ and get $\Gamma, \gamma \vdash_{\mathcal{S}} \beta$. By $\wedge$-introduction on the left,

$$\frac{\Gamma, \gamma \vdash_{\mathcal{S}} \beta}{\Gamma, \gamma \wedge \delta \vdash_{\mathcal{S}} \beta.}$$

But $\gamma \wedge \delta$, being an assumption, is already in Γ. Thus the conclusion is equivalent to $\Gamma \vdash_{\mathcal{S}} \beta$. $\square$

Obviously *the translation from $\mathcal{S}$ to $\mathcal{N}$ is not one-one*, since it was already not so for the implicational fragment alone. We now give a counterexample using $\wedge$ alone: the normal proof

$$\dfrac{\dfrac{\alpha \wedge \beta}{\alpha} \qquad \dfrac{\gamma \wedge \delta}{\gamma}}{\alpha \wedge \gamma}$$

is the translation of the following two proofs of the sequent

$$\alpha \wedge \beta, \gamma \wedge \delta \vdash_{\mathcal{S}} \alpha \wedge \gamma,$$

originated by two different orders of $\wedge$-introduction on the left:

$$\begin{array}{c} \alpha \vdash_{\mathcal{S}} \alpha \qquad \gamma \vdash_{\mathcal{S}} \gamma \\ \hline \alpha, \gamma \vdash_{\mathcal{S}} \alpha \wedge \gamma \\ \hline \alpha \wedge \beta, \gamma \vdash_{\mathcal{S}} \alpha \wedge \gamma \\ \hline \alpha \wedge \beta, \gamma \wedge \delta \vdash_{\mathcal{S}} \alpha \wedge \gamma \end{array} \qquad \text{and} \qquad \begin{array}{c} \alpha \vdash_{\mathcal{S}} \alpha \qquad \gamma \vdash_{\mathcal{S}} \gamma \\ \hline \alpha, \gamma \vdash_{\mathcal{S}} \alpha \wedge \gamma \\ \hline \alpha, \gamma \wedge \delta \vdash_{\mathcal{S}} \alpha \wedge \gamma \\ \hline \alpha \wedge \beta, \gamma \wedge \delta \vdash_{\mathcal{S}} \alpha \wedge \gamma. \end{array}$$

4.2 Semantics

Beth-Kripke models

The notions of *intuitionistic possible world* (or *Beth-Kripke model*) and *intuitionistic logical consequence* do not refer to connectives, and can thus be retained in their original forms 2.2.1 and 2.2.4. What needs to be supplemented is the definition of forcing 2.2.2.

Definition 4.2.1 Forcing (Cohen [1963], Kripke [1963]) *For a given possible world $\mathcal{A}$, the relation $\Vdash_{\mathcal{A}}$ defined in 2.2.2 is extended to conjunction as follows:*

$$\sigma \Vdash_{\mathcal{A}} \alpha \wedge \beta \quad \Leftrightarrow \quad \sigma \Vdash_{\mathcal{A}} \alpha \ \textit{ and } \ \sigma \Vdash_{\mathcal{A}} \beta.$$

The next result shows that the extension of forcing to $\wedge$ captures the intended meaning of conjunction.

Theorem 4.2.2 Intuitionistic Soundness and Completeness (Beth [1956], Kripke [1963]) *For any Γ and α:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \ \Leftrightarrow \ \Gamma \models_i \alpha.$$

Proof. For the Soundness direction, we supplement the proof of 2.2.5 by the cases dealing with conjunction.

- If $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \beta$ by $\wedge$-introduction, then $\Gamma \models_i \alpha$ and $\Gamma \models_i \beta$ by the induction hypothesis. Let σ be any state that forces all formulas of Γ in some world $\mathcal{A}$: then σ forces α and β, and hence it forces $\alpha \wedge \beta$ by definition of forcing. Since σ and $\mathcal{A}$ are arbitrary, $\Gamma \models_i \alpha \wedge \beta$.

- If $\Gamma \vdash_{\mathcal{N}} \alpha$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ by $\wedge$-elimination, then $\Gamma \models_i \alpha \wedge \beta$ by the induction hypothesis. Let σ be any state that forces all formulas of Γ in some world $\mathcal{A}$. By the induction hypothesis, σ forces $\alpha \wedge \beta$, and hence it forces α by definition of forcing. Since σ and $\mathcal{A}$ are arbitrary, $\Gamma \models_i \alpha$. Similarly for the other $\wedge$-elimination rule.

For the Completeness direction, we supplement the proof of 2.2.6 (to which we refer) by the case of conjunction. We have to prove that

$$\Theta \Vdash_{\mathcal{A}} \gamma \wedge \delta \ \Leftrightarrow \ \gamma \wedge \delta \in \Theta,$$

where Θ is any set of formulas closed under $\vdash_{\mathcal{N}}$.

- Suppose $\Theta \Vdash_{\mathcal{A}} \gamma \wedge \delta$. Then $\Theta \Vdash_{\mathcal{A}} \gamma$ and $\Theta \Vdash_{\mathcal{A}} \delta$, and by the induction hypothesis $\gamma \in \Theta$ and $\delta \in \Theta$. We want $\gamma \wedge \delta \in \Theta$. Suppose $\gamma \wedge \delta \notin \Theta$. By closure under $\vdash_{\mathcal{N}}$, $\Theta \nvdash_{\mathcal{N}} \gamma \wedge \delta$. By $\wedge$- *introduction*, $\Theta \nvdash_{\mathcal{N}} \gamma$ or $\Theta \nvdash_{\mathcal{N}} \delta$, contradiction (because Θ contains both γ and δ).

- Suppose $\gamma \wedge \delta \in \Theta$. We want $\Theta \Vdash_{\mathcal{A}} \gamma \wedge \delta$, i.e. $\Theta \Vdash_{\mathcal{A}} \gamma$ and $\Theta \Vdash_{\mathcal{A}} \delta$ or, by the induction hypothesis, $\gamma \in \Theta$ and $\delta \in \Theta$. But if $\gamma \wedge \delta$ is in Θ, then so are both γ and δ, by $\wedge$-*elimination* and closure under $\vdash_{\mathcal{N}}$. □

In particular, we get a *semantical proof of the Cut Elimination Theorem*, as in 2.2.7.

Exercises 4.2.3 a) *The Countable Model Property continues to hold.* (Hint: see 2.2.9.)
b) *The Constructive Model Property continues to hold.* (Hint: see 2.2.10.)
c) *The Finite Model Property continues to hold.* (Hint: see 2.2.11.)

Intuitionistic tableaux

The notion of *provability by intuitionistic tableaux* does not refer to connectives, and can thus be retained in the original form 2.3.2. What needs to be supplemented is the definition of tableaux 2.3.1.

Definition 4.2.4 *An* **intuitionistic tableau** *is a tree with nodes consisting of signed forcing assertions of the form $T\sigma \Vdash \alpha$ or $F\sigma \Vdash \alpha$, and consistent with the formation rules of 2.3.1, as well as with the following:*

3. *If a node $T\sigma \Vdash \alpha \wedge \beta$ is on the tree, then we can extend any branch going through it by adding $T\sigma \Vdash \alpha$ and $T\sigma \Vdash \beta$. Graphically,*

$$\begin{array}{c} T\sigma \Vdash \alpha \wedge \beta \\ \hline\hline T\sigma \Vdash \alpha \\ \hline T\sigma \Vdash \beta, \end{array}$$

where the double line shows that the bottom nodes do not have to immediately follow the top one.

4. *If a node $F\sigma \Vdash \alpha \wedge \beta$ is on the tree, then we can split any branch going through it by adding $F\sigma \Vdash \alpha$ in one direction and $F\sigma \Vdash \beta$ in the other. Graphically,*

$$\begin{array}{c} F\sigma \Vdash \alpha \wedge \beta \\ \hline\hline F\sigma \Vdash \alpha \qquad F\sigma \Vdash \beta. \end{array}$$

The next result shows that the extension of the tableaux rules to $\wedge$ captures the intended meaning of validity.

Theorem 4.2.5 Soundness and Completeness for Tableaux (Nerode [1990]) *For any Γ and α:*

$$\Gamma \vdash_{\mathcal{T}} \alpha \;\Leftrightarrow\; \Gamma \models_i \alpha.$$

Proof. The proofs of 2.3.4 and 2.3.5 are easily supplemented by the cases dealing with conjunction, since the rules for the construction of tableaux were chosen to mirror the definition of forcing. □

In conclusion, we have proved that all systems considered so far are equivalent: for any Γ and α,

$$\Gamma \vdash_{\mathcal{N}} \alpha \;\Leftrightarrow\; \Gamma \vdash_{\mathcal{H}} \alpha \;\Leftrightarrow\; \Gamma \vdash_{\mathcal{S}} \alpha \;\Leftrightarrow\; \Gamma \vdash_{\mathcal{T}} \alpha \;\Leftrightarrow\; \Gamma \models_i \alpha.$$

4.3 Complexity

æ

Part B

Categories

Chapter 5

Heyting $\sqcap$-Algebras

The gaol of the present chapter is to introduce an approach to semantics different from (but, as it will turn out, related to) the one via Beth-Kripke models. The main idea is to exploit the extra power provided by conjunction, which allows to *identify finite sets of premises with their conjunction.* This is expressed by the following equivalence:

$$\gamma_1, \ldots, \gamma_n \vdash_{\mathcal{N}} \alpha \;\Leftrightarrow\; \gamma_1 \wedge \cdots \wedge \gamma_n \vdash_{\mathcal{N}} \alpha,$$

which holds by associativity of ∧, and successive applications of the ∧-elimination and ∧-introduction rules.

By an application of →-introduction we also get the following *strengthening of the Deduction Theorem*:

$$\frac{\gamma_1, \ldots, \gamma_n \vdash_{\mathcal{N}} \alpha}{\vdash_{\mathcal{N}} \gamma_1 \wedge \cdots \wedge \gamma_n \to \alpha.}$$

This shows that there is a strong interplay among the relation $\vdash_{\mathcal{N}}$ and the connectives ∧ and →, which we now attempt to describe in a purely algebraic way.

We will provide the needed algebraic background, and refer to Grätzer [1978], and Davey and Priestley [1990] for more detailed treatments. Similarly, Rasiowa and Sikorski [1963], and Rasiowa [1974] are the references for detailed treatments of the connections with logic.

5.1 Heyting $\sqcap$-Algebras

Algebraic models

To describe the interplay among $\vdash_{\mathcal{N}}$, ∧ and → we consider algebraic structures $\mathcal{A}$ with an underlying set A, one relation $\sqsubseteq_{\mathcal{A}}$ on A intended to model $\vdash_{\mathcal{N}}$, and two binary operations $\sqcap_{\mathcal{A}}$ and $\Rightarrow_{\mathcal{A}}$ on A intended to model $\wedge$ and $\rightarrow$, respectively. Such a structure will be used to define an interpretation function $[\![\ ]\!]^{\mathcal{A}}_{\rho}$ for formulas, relative to a given interpretation (called *environment*) of the propositional letters, i.e. to a function ρ from the set of all propositional letters to the underlying set A of $\mathcal{A}$.

As usual, we will drop the superscript $\mathcal{A}$ or the subscript ρ when no confusion arises. In particular, there should be no confusion about the uses of $\Rightarrow$ as a metamathematical sign, i.e. as an abbreviation for an informal implication, and as an algebraic sign, i.e. as an interpretation of a formal implication.

Definition 5.1.1 Canonical Algebraic Interpretation. *Given a structure*

$$\mathcal{A} = \langle A, \sqcap, \Rightarrow \rangle$$

and an **environment $\boldsymbol{\rho}$** *on it, i.e. a function*

$$\rho : \textit{Propositional Letters} \longrightarrow A,$$

we define the **canonical algebraic interpretation $[\![\ ]\!]_{\boldsymbol{\rho}}$** *by induction on formulas, as follows:*

$$[\![\alpha]\!]_{\rho} = \begin{cases} \rho(p) & \textit{if } \alpha = p \\ [\![\beta]\!]_{\rho} \sqcap [\![\gamma]\!]_{\rho} & \textit{if } \alpha = \beta \wedge \gamma \\ [\![\beta]\!]_{\rho} \Rightarrow [\![\gamma]\!]_{\rho} & \textit{if } \alpha = \beta \rightarrow \gamma. \end{cases}$$

By induction, $[\![\alpha]\!]_{\rho} \in A$ for every α. Having the notion of interpretation, we define the notion of model by mirroring $\vdash_{\mathcal{N}}$, modulo the identification of finite sets of premises with their conjunctions.

Definition 5.1.2 *An* **algebraic model** *of Implicational Calculus with Conjunction is a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, \sqcap, \Rightarrow \rangle$$

such that, for every $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ and α,

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies (\forall \rho)([\![\gamma_1]\!]_{\rho} \sqcap \cdots \sqcap [\![\gamma_n]\!]_{\rho} \sqsubseteq [\![\alpha]\!]_{\rho}).$$

Motivation

Our next goal is to determine conditions on $\mathcal{A}$ ensuring that $\mathcal{A}$ is an algebraic model. The following observations on $\mathcal{N}$ provide sufficient conditions.

1. *provable equivalence defines an equivalence relation on formulas, and $\vdash_{\mathcal{N}}$ induces a partial ordering on the equivalence classes*
   Axioms provide *reflexivity*:
   $$\alpha \vdash_{\mathcal{N}} \alpha.$$

The possibility of merging proofs as follows:

$$\begin{array}{c} \alpha \\ \mathcal{D}_\beta \\ \beta \end{array} \quad \text{and} \quad \begin{array}{c} \beta \\ \mathcal{D}_\gamma \\ \gamma \end{array} \quad \text{into} \quad \begin{array}{c} \alpha \\ \mathcal{D}_\beta \\ \beta \\ \mathcal{D}_\gamma \\ \gamma \end{array}$$

provides *transitivity*:

$$(\alpha \vdash_{\mathcal{N}} \beta) \wedge (\beta \vdash_{\mathcal{N}} \gamma) \implies (\alpha \vdash_{\mathcal{N}} \gamma).$$

The following deduction:

$$\dfrac{\dfrac{\dfrac{\alpha \vdash_{\mathcal{N}} \beta}{\vdash_{\mathcal{N}} \alpha \to \beta} \qquad \dfrac{\beta \vdash_{\mathcal{N}} \alpha}{\vdash_{\mathcal{N}} \beta \to \alpha}}{\vdash_{\mathcal{N}} (\alpha \to \beta) \wedge (\beta \to \alpha)}}{\vdash_{\mathcal{N}} \alpha \leftrightarrow \beta}$$

provides *antisymmetry*:

$$(\alpha \vdash_{\mathcal{N}} \beta) \wedge (\beta \vdash_{\mathcal{N}} \alpha) \implies (\vdash_{\mathcal{N}} \alpha \leftrightarrow \beta).$$

In the following we will abuse language, and use $\vdash_{\mathcal{N}}$ to refer to the associated partial ordering on the equivalence classes.

2. $\vdash_{\mathcal{N}}$ *admits a greatest element*
   If β is any formula provable without any assumption, i.e. such that $\vdash_{\mathcal{N}} \beta$, then by the Thinning Rule we also have $\alpha \vdash_{\mathcal{N}} \beta$ for any formula α, and thus the equivalence class of β is the greatest element w.r.t. $\vdash_{\mathcal{N}}$.

3. $\wedge$ *induces the greatest lower bound operation on the equivalence classes*
   The $\wedge$-elimination rules show that $\wedge$ induces a *lower bound*:

$$\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta}{\Gamma \vdash_{\mathcal{N}} \alpha}$$

   says that anything less than $\alpha \wedge \beta$ must be less than α. Similarly for β.

   The $\wedge$-introduction rule shows that $\wedge$ induces the *greatest* lower bound:

$$\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta}$$

   says that anything below both α and β must also be below $\alpha \wedge \beta$.

We leave to the reader the trivial check that the operation induced by $\wedge$ is well-defined on equivalence classes, in the sense that

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \leftrightarrow \alpha' \qquad \Gamma \vdash_{\mathcal{N}} \beta \leftrightarrow \beta'}{\Gamma \vdash_{\mathcal{N}} (\alpha \wedge \beta) \leftrightarrow (\alpha' \wedge \beta').}$$

4. *$\to$ induces an operation on the equivalence classes that behaves on the right of $\vdash_{\mathcal{N}}$ as the operation induced by $\wedge$ does on the left*
Precisely,

$$(\alpha \wedge \beta) \vdash_{\mathcal{N}} \gamma \;\Leftrightarrow\; \alpha \vdash_{\mathcal{N}} (\beta \to \gamma).$$

This is just a symmetric restatement of the $\to$-introduction and $\to$-elimination rules, modulo the additional freedom allowed by $\wedge$ in the treatment of the hypotheses.

For the left to right direction, given a proof $\mathcal{D}$ of γ from $\alpha \wedge \beta$, by $\wedge$-introduction and $\to$-introduction we get:

$$\begin{array}{c} \dfrac{\alpha \qquad [\beta]^{(1)}}{\alpha \wedge \beta} \\ \mathcal{D} \\ \dfrac{\gamma}{\beta^{(1)} \to \gamma.} \end{array}$$

For the right to left direction, given a proof $\mathcal{D}$ of $\beta \to \gamma$ from α, by $\wedge$-elimination and $\to$-elimination we get:

$$\dfrac{\dfrac{\alpha \wedge \beta}{\beta} \qquad \begin{array}{c} \dfrac{\alpha \wedge \beta}{\alpha} \\ \mathcal{D} \\ \beta \to \gamma \end{array}}{\gamma.}$$

We leave to the reader the trivial check that the operation induced by $\to$ is well-defined on equivalence classes, in the sense that

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \leftrightarrow \alpha' \qquad \Gamma \vdash_{\mathcal{N}} \beta \leftrightarrow \beta'}{\Gamma \vdash_{\mathcal{N}} (\alpha \to \beta) \leftrightarrow (\alpha' \to \beta').}$$

Adjointness

The situation depicted by Property 4 above can be abstracted as follows.

Definition 5.1.3 (Kan [1958]) *Given a partial ordering $\sqsubseteq$ on a set A, a function g on A is a* **right adjoint** *of a function f on A if, for every x and y in A,*

$$f(x) \sqsubseteq y \;\Leftrightarrow\; x \sqsubseteq g(y).$$

Then, by letting $f_\beta(x) = x \wedge \beta$ and $g_\beta(y) = \beta \to y$, condition 4 above says that g_β is a right adjoint of f_β w.r.t. $\vdash_{\mathcal{N}}$, for every formula β. We will abuse language and say that $\to$ *is a right adjoint of* $\wedge$ *w.r.t.* $\vdash_{\mathcal{N}}$.

The following simple fact shows that the adjointness condition is actually sufficient to uniquely determine the meaning of $\to$ in terms of $\vdash_{\mathcal{N}}$ and $\wedge$.

Proposition 5.1.4 Uniqueness of Right Adjoints. *Given a partial ordering $\sqsubseteq$ on a set A, a function f on A has at most one right adjoint w.r.t. $\sqsubseteq$.*

Proof. If g is any right adjoint of f, then

$$f(x) \sqsubseteq y \;\Leftrightarrow\; x \sqsubseteq g(y).$$

In particular, for $x = g(y)$,

$$f(g(y)) \sqsubseteq y \;\Leftrightarrow\; g(y) \sqsubseteq g(y).$$

Since the right-hand-side is always true, because $\sqsubseteq$ is a partial ordering, then

$$f(g(y)) \sqsubseteq y$$

for every y.

Let now g_1 and g_2 be two right adjoints of f. Then

$$f(x) \sqsubseteq y \;\Leftrightarrow\; x \sqsubseteq g_2(y)$$

because g_2 is a right adjoint. In particular, for $x = g_1(y)$,

$$f(g_1(y)) \sqsubseteq y \;\Leftrightarrow\; g_1(y) \sqsubseteq g_2(y).$$

Since the left-hand-side is always true, because g_1 is a right adjoint,

$$g_1(y) \sqsubseteq g_2(y).$$

By interchanging g_1 and g_2,

$$g_2(y) \sqsubseteq g_1(y).$$

Since $\sqsubseteq$ is a partial ordering,

$$g_1(y) = g_2(y).$$

Then $g_1 = g_2$. $\square$

Of course, there is nothing special about *right* adjoints. In definition 5.1.3 we say that f is a *left* adjoint of g, and a proof dual to the one above, using this time the fact that $x \sqsubseteq g(f(x))$, shows that left adjoints are unique when they exist.

The next result provides an alternative characterization of adjointness in terms of the properties discovered in the proof of uniqueness, and will be useful for generalizations in the next chapter.

Proposition 5.1.5 *If f and g are monotone functions on a partial ordering, then g is the right adjoint of f if and only if the following hold:*

1. $x \sqsubseteq g(f(x))$
2. $f(g(y)) \sqsubseteq y$.

Proof. The necessity of the condition has been proved in the proof of the previous result. For sufficiency, suppose first

$$f(x) \sqsubseteq y.$$

Then

$$x \sqsubseteq g(f(x)) \sqsubseteq g(y)$$

by property 1 and monotonicity of g. Conversely, suppose

$$x \sqsubseteq g(y).$$

Then

$$f(x) \sqsubseteq f(g(y)) \sqsubseteq y$$

by monotonicity of f and property 2. Thus

$$f(x) \sqsubseteq y \iff x \sqsubseteq g(y),$$

and g is the right adjoint of f. □

Properties 1 and 2 are called, respectively, the *unit* and *counit* of the adjointness. In the special case of the right adjoint of $\sqcap$, they translate into

$$x \sqsubseteq (y \Rightarrow (x \sqcap y)) \qquad \text{and} \qquad ((x \Rightarrow y) \sqcap x) \sqsubseteq y.$$

An improved presentation of adjointness in terms of equalities, instead of inequalities, will be given in 5.1.8.

Definition

Having discovered the algebraic properties forced on $\vdash_{\mathcal{N}}$, $\wedge$ and $\rightarrow$ by the logical rules, we now abstract them and introduce algebraic structures that, in the light of the previous discussion, turn out to be algebraic models.

Definition 5.1.6 (Ogasawara [1939], Birkhoff [1940], McKinsey and Tarski [1946]) *A* **Heyting $\sqcap$-algebra** *(read as '***inf-algebra***' or '***meet-algebra***') is a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, \Rightarrow, 1 \rangle$$

such that:

1. *$\sqsubseteq$ is a partial ordering with $=$ as associated equality*
2. *1 is the greatest element of A w.r.t. $\sqsubseteq$*
3. *$\sqcap$ is the g.l.b. operation associated with $\sqsubseteq$*
4. *$\Rightarrow$ is the right adjoint of $\sqcap$ w.r.t. $\sqsubseteq$.*

Conditions 1 and 3 define a **lowersemilattice**, i.e. a partially ordered structure in which every pair of elements has a g.l.b. This implies that every non empty finite subset has a g.l.b., but leaves open the degenerate case of the empty set, which is taken care of by condition 2 (since the g.l.b. of $\emptyset$ is the greatest element).

In other words, the first three conditions together establish that every finite subset, empty or not, has a g.l.b.

Equational presentation ⋆

The definition of a Heyting $\sqcap$-algebra can be rephrased in various ways. Here we are interested in purely equational presentations, which show that Heyting $\sqcap$-algebras are *algebraic varieties*.

Knowing that $\sqsubseteq$ and $\sqcap$ must be related as said, we can take either one of them as primitive and define the other one accordingly. Precisely, we can either ask $\sqsubseteq$ to be a partial ordering and $\sqcap$ to be its associated g.l.b., or we can impose conditions on $\sqcap$ that would ensure that there is a unique partial ordering on A such that $\sqcap$ is its associated g.l.b., and define $\sqsubseteq$ to be such an ordering.

Proposition 5.1.7 Equational Presentation of Lowersemilattices with Greatest Element (Huntington [1904]) *In a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, 1 \rangle$$

$\sqsubseteq$ is a partial ordering with $=$ as associated equality, 1 as greatest element and $\sqcap$ as associated g.l.b. if and only if

$$x \sqsubseteq y \;\Leftrightarrow\; (x \sqcap y) = x,$$

and the following hold:

1. *$x \sqcap x = x$ (**idempotency**)*
2. *$x \sqcap (y \sqcap z) = (x \sqcap y) \sqcap z$ (**associativity**)*
3. *$x \sqcap y = y \sqcap x$ (**commutativity**)*
4. *$x \sqcap 1 = x$*

Proof. The stated properties are obviously necessary. Conversely, suppose they hold: we show that $\sqsubseteq$, defined as

$$x \sqsubseteq y \;\Leftrightarrow\; (x \sqcap y) = x,$$

is a partial ordering.

1. *reflexivity*

   $x \sqsubseteq x$ is translated as $x \sqcap x = x$, which holds by property 1.

2. *transitivity*

   If $x \sqsubseteq y$ and $y \sqsubseteq z$, then $x \sqcap y = x$ and $y \sqcap z = y$. By property 2

   $$x \sqcap z = (x \sqcap y) \sqcap z = x \sqcap (y \sqcap z) = x \sqcap y = x,$$

   i.e. $x \sqsubseteq z$.

3. *antisymmetry*

   If $x \sqsubseteq y$ and $y \sqsubseteq x$, then $x \sqcap y = x$ and $y \sqcap x = y$. By property 3, $x = y$.

4. *greatest element*

   By property 4, $x \sqcap 1 = x$, i.e. $x \sqsubseteq 1$ for any x. □

We turn now to the additional adjointness condition.

Proposition 5.1.8 Equational Presentation of Adjointness (Monteiro [1955], Rasiowa and Sikorski [1963]) *In a lowersemilattice*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, \Rightarrow, 1 \rangle$$

(where $\sqsubseteq$ is a partial ordering, $=$ the associated equality, $\sqcap$ the associated g.l.b. and 1 the greatest element) $\Rightarrow$ is the right adjoint of $\sqcap$ w.r.t. $\sqsubseteq$ if and only if the following hold:

5. $(x \Rightarrow x) = 1$

6. $x \Rightarrow (y \sqcap z) = (x \Rightarrow y) \sqcap (x \Rightarrow z)$

7. $x \sqcap (x \Rightarrow y) = x \sqcap y$

8. $y \sqcap (x \Rightarrow y) = y$.

Proof. We first prove that the stated properties are necessary.

- $(x \Rightarrow x) = 1$

  It is enough to show that
  $$a \sqsubseteq (x \Rightarrow x)$$
  holds for every a. But this means
  $$a \sqcap x \sqsubseteq x,$$
  which is always true by definition of $\sqcap$.

- $x \Rightarrow (y \sqcap z) = (x \Rightarrow y) \sqcap (x \Rightarrow z)$

  It is enough to note the following equivalences:
  $$\begin{array}{ll} & a \sqsubseteq x \Rightarrow (y \sqcap z) \\ \Longleftrightarrow & a \sqcap x \sqsubseteq y \sqcap z \\ \Longleftrightarrow & (a \sqcap x \sqsubseteq y) \ \wedge \ (a \sqcap x \sqsubseteq z) \\ \Longleftrightarrow & (a \sqsubseteq x \Rightarrow y) \ \wedge \ (a \sqsubseteq x \Rightarrow z) \\ \Longleftrightarrow & a \sqsubseteq (x \Rightarrow y) \sqcap (x \Rightarrow z), \end{array}$$
  which hold by adjointness and definition of $\sqcap$ (twice).

- $x \sqcap (x \Rightarrow y) = x \sqcap y$

  It is enough to note the following equivalences:
  $$\begin{array}{ll} & a \sqsubseteq x \sqcap (x \Rightarrow y) \\ \Longleftrightarrow & (a \sqsubseteq x) \ \wedge \ (a \sqsubseteq x \Rightarrow y) \\ \Longleftrightarrow & (a \sqsubseteq x) \ \wedge \ (a \sqcap x \sqsubseteq y) \\ \Longleftrightarrow & (a \sqsubseteq x) \ \wedge \ (a \sqsubseteq y) \\ \Longleftrightarrow & a \sqsubseteq x \sqcap y, \end{array}$$
  which hold by definition of $\sqcap$ and adjointness.

- $y \sqcap (x \Rightarrow y) = y$

  By definition of $\sqcap$ this is equivalent to

  $$y \sqsubseteq (x \Rightarrow y),$$

  and by adjointness it is equivalent to

  $$y \sqcap x \sqsubseteq y,$$

  which is obviously true.

We now prove that the stated properties are sufficient, by showing that if $\Rightarrow$ satifies them, then

$$(x \sqcap y) \sqsubseteq z \iff x \sqsubseteq (y \Rightarrow z).$$

The right to left direction is immediate. If

$$x \sqsubseteq (y \Rightarrow z),$$

then

$$x \sqcap y \sqsubseteq y \sqcap (y \Rightarrow z),$$

and by 7

$$x \sqcap y \sqsubseteq y \sqcap z \sqsubseteq z.$$

For the left to right direction, suppose

$$(x \sqcap y) \sqsubseteq z.$$

If we prove that

$$[y \Rightarrow (x \sqcap y)] \sqsubseteq (y \Rightarrow z),$$

then we have: by 6,

$$(y \Rightarrow x) \sqcap (y \Rightarrow y) \sqsubseteq (y \Rightarrow z);$$

by 5,

$$(y \Rightarrow x) \sqsubseteq (y \Rightarrow z);$$

and by 8, which implies $x \sqsubseteq (y \Rightarrow x)$,

$$x \sqsubseteq (y \Rightarrow x) \sqsubseteq (y \Rightarrow z).$$

It remains to prove the following *monotonicity property* (to be applied above with $b = x \sqcap y$ and $c = z$):

$$\text{if } \ b \sqsubseteq c, \ \text{ then } \ (a \Rightarrow b) \sqsubseteq (a \Rightarrow c).$$

But if $b \sqsubseteq c$, then $b \sqcap c = b$. So

$$a \Rightarrow (b \sqcap c) = (a \Rightarrow b),$$

and by 6

$$(a \Rightarrow b) \sqcap (a \Rightarrow c) = (a \Rightarrow b),$$

i.e.

$$(a \Rightarrow b) \sqsubseteq (a \Rightarrow c). \quad \Box$$

Exercises 5.1.9 a) $y \sqsubseteq (x \Rightarrow y)$. (Hint: from 8.)
b) $x \sqsubseteq (x \Rightarrow y)$ *if and only if* $x \sqsubseteq y$. (Hint: from 7.)
c) $(1 \Rightarrow x) = x$. (Hint: from 7.)
d) $(x \Rightarrow y) = 1$ *if and only if* $x \sqsubseteq y$. (Hint: from 7.)

Filters and quotients ⋆

The proof of the Intuitionistic Completeness Theorem 2.2.6 and 4.2.2 shows that we can model the behavior of $\vdash_{\mathcal{N}}$, $\wedge$ and $\rightarrow$ by considering sets of formulas closed under $\vdash_{\mathcal{N}}$ and $\wedge$. It is thus natural to try to adapt the same proof, and to model the behavior of $\sqsubseteq$, $\sqcap$ and $\Rightarrow$ by considering sets of elements with the same closure properties.

Definition 5.1.10 (Cartan [1937]) *Given a Heyting* $\sqcap$*-algebra*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, \Rightarrow, 1 \rangle,$$

a nonempty subset F *of* A *is a* **filter** *if it is:*

- *upward closed under* $\sqsubseteq$, *i.e. for any* x *and* y *in* A,

$$x \in F \wedge x \sqsubseteq y \implies y \in F$$

- *closed under* $\sqcap$, *i.e. for any* x *and* y *in* A,

$$x \in F \wedge y \in F \implies x \sqcap y \in F.$$

The definition of filter was given in terms of $\sqsubseteq$ and $\wedge$, but the next result shows that we could have used $\Rightarrow$ instead.

Proposition 5.1.11 *In a Heyting* $\sqcap$*-algebra, a subset* F *is a filter if and only if:*

1. $1 \in F$
2. *if* $a \in F$ *and* $a \Rightarrow b \in F$ *then* $b \in F$.

Proof. The conditions are obviously necessary. A filter contains at least an element a, because it is a nonempty subset, and hence it contains 1 by upward closure, since $a \sqsubseteq 1$. And if a filter contains a and $a \Rightarrow b$, then it contains $a \sqcap (a \Rightarrow b)$ by closure under $\sqcap$, $a \sqcap b$ by Property 7 of 5.1.8, and hence b by upward closure, since $a \sqcap b \sqsubseteq b$.

We now show that the two conditions are sufficient:

- *if $a, b \in F$, then $a \sqcap b \in F$*

  By 2 (twice), it is enough to show that $a \Rightarrow (b \Rightarrow a \sqcap b) \in F$. But

  $$\begin{aligned} a \sqcap b \sqsubseteq a \sqcap b &\iff a \sqsubseteq (b \Rightarrow a \sqcap b) \\ &\iff (a \Rightarrow a) \sqsubseteq [a \Rightarrow (b \Rightarrow a \sqcap b)] \\ &\iff 1 \sqsubseteq [a \Rightarrow (b \Rightarrow a \sqcap b)] \end{aligned}$$

  by adjointness, monotonicity, and Property 5 of 5.1.8.

  Thus $[a \Rightarrow (b \Rightarrow a \sqcap b)] = 1 \in F$ by 1.

- *if $a \in F$ and $a \sqsubseteq b$, then $b \in F$*

  If $a \sqsubseteq b$, then $(a \Rightarrow a) \sqsubseteq (a \Rightarrow b)$ by monotonicity, and $1 \sqsubseteq (a \Rightarrow b)$ by Property 5 of 5.1.8.

  Now $(a \Rightarrow b) = 1 \in F$ by 1, and then $b \in F$ by 2. □

Exercises 5.1.12 Filter generated by a set. Given a Heyting $\sqcap$-algebra A and a subset B of it, the **filter generated by B** is the smallest filter on A containing B.

a) *The filter generated by B exists.* (Hint: the intersection of a non empty family of filters is a filter, and the family of filters containing B is not empty, because the set A is in it.)

b) *The filter generated by B is the upward closure of the set of all finite meets of elements of B.* (Hint: the upward closure of the set of all finite meets of elements of B is a filter, and is contained in every filter containing B.)

Part of the interest of the notion of filter lies in the fact that it can be used for the definition of quotient algebras.

Proposition 5.1.13 *Given a Heyting $\sqcap$-algebra $\mathcal{A}$ and a filter F on it, the relation*

$$x \sqsubseteq_F y \iff (x \Rightarrow y) \in F$$

induces an equivalence relation, whose set of equivalence classes

$$\mathcal{A}_{/F} = \{[x] : x \in A\}$$

is a Heyting $\sqcap$-algebra (called the **quotient of $\boldsymbol{A}$ w.r.t. $\boldsymbol{F}$***), with operations and relations induced from those of $\mathcal{A}$.*

Proof. We prove the following facts:

- *the relation $\sqsubseteq_F$ induces an equivalence relation*

  It is enough to show that $\sqsubseteq_F$ is reflexive and transitive, since then the relation

  $$x =_F y \iff x \sqsubseteq_F y \wedge y \sqsubseteq_F x$$

  is reflexive, transitive and symmetric, i.e. an equivalence relation.

  Reflexivity $x \sqsubseteq_F x$ follows from Property 5 of 5.1.8, since $(x \Rightarrow x) = 1 \in F$.

  For transitivity, suppose

  $$x \sqsubseteq_F y \qquad \text{and} \qquad y \sqsubseteq_F z,$$

  i.e.

  $$(x \Rightarrow y) \in F \qquad \text{and} \qquad (y \Rightarrow z) \in F.$$

  By closure of F under $\sqcap$, $(x \Rightarrow y) \sqcap (y \Rightarrow z) \in F$. If we prove that

  $$(x \Rightarrow y) \sqcap (y \Rightarrow z) \sqsubseteq (x \Rightarrow z),$$

  then $(x \Rightarrow z) \in F$, i.e. $x \sqsubseteq_F z$, follows by upward closure of F.

  Now

  $$\begin{aligned} & (x \Rightarrow y) \sqcap (y \Rightarrow z) \sqsubseteq (x \Rightarrow z) \\ \iff\ & x \sqcap (x \Rightarrow y) \sqcap (y \Rightarrow z) \sqsubseteq z \\ \iff\ & x \sqcap y \sqcap (y \Rightarrow z) \sqsubseteq z \\ \iff\ & x \sqcap y \sqcap z \sqsubseteq z \end{aligned}$$

  by adjointness and Property 7 of 5.1.8 twice, and the last line is obviously true.

- *the operation $\sqcap$ induces the g.l.b.*

  $\sqcap$ induces on $\mathcal{A}_{/F}$ the operation

  $$[x] \sqcap_F [y] = [x \sqcap y].$$

  To prove that this gives a *lower bound*, we need to show that $(x \sqcap y) \sqsubseteq_F x$, i.e. that $(x \sqcap y \Rightarrow x) \in F$. If we prove that

  $$(x \Rightarrow x) \sqsubseteq (x \sqcap y \Rightarrow x),$$

  then we get $(x \sqcap y \Rightarrow x) \in F$ from the fact that $(x \Rightarrow x) = 1$ by Property 5 of 5.1.8, since a filter contains 1 and is closed upwards.

It remains to prove the following *contravariance property* (to be applied above with $a = x \sqcap y$, $b = x$ and $c = x$):

$$\text{if } \ a \sqsubseteq b, \ \text{ then } \ (b \Rightarrow c) \sqsubseteq (a \Rightarrow c).$$

For this it is enough to note that if

$$x \sqsubseteq (b \Rightarrow c),$$

then

$$x \sqcap b \sqsubseteq c$$

by adjointness,

$$x \sqcap a \sqsubseteq c$$

because $a \sqsubseteq b$, and

$$x \sqsubseteq (a \Rightarrow c)$$

by adjointness.

To prove that $\sqcap$ induces the *greatest* lower bound, suppose

$$z \sqsubseteq_F x \qquad \text{and} \qquad z \sqsubseteq_F y,$$

i.e.

$$(z \Rightarrow x) \in F \qquad \text{and} \qquad (z \Rightarrow y) \in F.$$

Then

$$(z \Rightarrow x) \sqcap (z \Rightarrow y) \in F$$

by closure under $\sqcap$, and

$$z \Rightarrow (x \sqcap y) \in F$$

by Property 6 of 5.1.8. Thus

$$z \sqsubseteq_F (x \sqcap y).$$

- *the operation* $\Rightarrow$ *induces a right adjoint*

  $\Rightarrow$ induces on $\mathcal{A}_{/F}$ the operation

$$[x] \Rightarrow_F [y] = [x \Rightarrow y].$$

To prove that this gives a right adjoint of $\sqcap_F$, we need to show that

$$[x] \sqcap_F [y] \sqsubseteq_F [z] \iff [x] \sqsubseteq_F ([y] \Rightarrow_F [z]),$$

i.e.

$$[x \sqcap y] \sqsubseteq_F [z] \iff [x] \sqsubseteq_F [y \Rightarrow z],$$

and hence

$$(x \sqcap y \Rightarrow z) \in F \iff x \Rightarrow (y \Rightarrow z) \in F.$$

This follows from

$$\begin{array}{rl} & a \sqsubseteq (x \sqcap y) \Rightarrow z \\ \iff & a \sqcap x \sqcap y \sqsubseteq z \\ \iff & a \sqcap x \sqsubseteq (y \Rightarrow z) \\ \iff & a \sqsubseteq x \Rightarrow (y \Rightarrow z), \end{array}$$

by adjointness.

- *the class of equivalence of* 1 *is the greatest element*

  It is enough to show that $x \sqsubseteq_F 1$, i.e. $(x \Rightarrow 1) \in F$, for any x. But $1 \sqsubseteq (x \Rightarrow 1)$ by Property 8 of 5.1.8, and hence $(x \Rightarrow 1) = 1 \in F$. □

Exercise 5.1.14 *If F is a filter, then $F = \{a : a =_F 1\}$, i.e.* $[1] = F$. (Hint: it is enough to show that $x \in F$ if and only if $1 \sqsubseteq_F x$. But a filter contains 1 and is closed under $\Rightarrow$. Thus, if $1 \Rightarrow x$ is in F, so is x. Conversely, from $x \sqsubseteq (1 \Rightarrow x)$ it follows that if x is in F, then so is $1 \Rightarrow x$.)

In applications, a typical way of building a quotient of a Heyting $\sqcap$-algebra A is via an onto homomorphism $f : A \to B$. Then

$$F = \{x : f(x) = 1\}$$

is a filter, for the following reasons:

- F is upward closed, since if $x \sqsubseteq y$, then $f(x) \sqsubseteq f(y)$ because f preserves $\sqsubseteq$, and thus if $f(x) = 1$, then $f(y) = 1$.

- F is closed under $\sqcap$, since $f(x \sqcap y) = f(x) \sqcap f(y)$ because f preserves $\sqcap$, and thus if $f(x) = f(y) = 1$, then $f(x \sqcap y) = 1$.

Moreover,

$$(x \Rightarrow y) \in F \iff f(x \Rightarrow y) = 1 \iff (f(x) \Rightarrow f(y)) = 1 \iff f(x) \sqsubseteq f(y),$$

i.e.

$$x =_F y \iff f(x) = f(y).$$

In other words, we identify classes on A with elements of the range of f. If f is onto, then the range of f is B, and $A_{/F}$ is isomorphic to B.

5.2 Soundness and Completeness Theorem

We can view Heyting $\sqcap$-algebras as an analogue of intuitionistic worlds. *Local truth* now means being evaluated to 1 under a given environment, and *global truth* means being evaluated to 1 under every environment. The following notion is then the analogue of intuitionistic logical validity (2.2.4).

Definition 5.2.1 *A formula α is an* **algebraic consequence** *of $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ (written $\boldsymbol{\Gamma \models_a \alpha}$) if for every Heyting $\sqcap$-algebra $\mathcal{A}$ and every environment ρ on it,*

$$(\llbracket \gamma_1 \rrbracket_\rho \sqcap \cdots \sqcap \llbracket \gamma_n \rrbracket_\rho) \sqsubseteq \llbracket \alpha \rrbracket_\rho .$$

In the limit case of Γ empty, we get the notion of algebraic validity: α is **algebraically valid** (written $\models_a \alpha$) if α evaluates to 1 in every Heyting $\sqcap$-algebra, under every environment.

Lindenbaum algebras

We first show that any Heyting $\sqcap$-algebra is a model of the Intuitionistic Implicational Calculus with Conjunction.

Theorem 5.2.2 Algebraic Soundness (McKinsey and Tarski [1948], Rasiowa [1951]) *For any Γ and α,*

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies \Gamma \models_a \alpha .$$

Proof. We prove that if $\mathcal{A}$ is a Heyting $\sqcap$-algebra, then it is an algebraic model, i.e. if ρ is an environment on $\mathcal{A}$ and $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ and α are given, then

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies (\llbracket \gamma_1 \rrbracket_\rho \sqcap \cdots \sqcap \llbracket \gamma_n \rrbracket_\rho) \sqsubseteq \llbracket \alpha \rrbracket_\rho .$$

The definition of a Heyting $\sqcap$-algebra has been discovered precisely by looking at the properties that would make the present proof work, and thus we only have to repeat the work done above, proceeding by induction on $\vdash_{\mathcal{N}}$. For simplicity we write $\llbracket \Gamma \rrbracket_\rho$ for $\llbracket \gamma_1 \rrbracket_\rho \sqcap \cdots \sqcap \llbracket \gamma_n \rrbracket_\rho$.

If $\Gamma, \beta \vdash_{\mathcal{N}} \beta$ is an assumption, then

$$(\llbracket \Gamma \rrbracket_\rho \sqcap \llbracket \beta \rrbracket_\rho) \sqsubseteq \llbracket \beta \rrbracket_\rho$$

follows from the fact that $\sqcap$ is a lower bound w.r.t. to $\sqsubseteq$.

If $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ is obtained from $\Gamma, \alpha \vdash_{\mathcal{N}} \beta$ by $\to$-introduction, then

$$(\llbracket \Gamma \rrbracket_\rho \sqcap \llbracket \alpha \rrbracket_\rho) \sqsubseteq \llbracket \beta \rrbracket_\rho$$

by the induction hypothesis, and hence

$$[\![\Gamma]\!]_\rho \sqsubseteq ([\![\alpha]\!]_\rho \Rightarrow [\![\beta]\!]_\rho) = [\![\alpha \to \beta]\!]_\rho$$

by right adjointness of $\Rightarrow$ and definition of $[\![\ ]\!]_\rho$.

If $\Gamma \vdash_{\mathcal{N}} \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ by →-elimination, then

$$[\![\Gamma]\!]_\rho \sqsubseteq [\![\alpha]\!]_\rho \qquad \text{and} \qquad [\![\Gamma]\!]_\rho \sqsubseteq [\![\alpha \to \beta]\!]_\rho = ([\![\alpha]\!]_\rho \Rightarrow [\![\beta]\!]_\rho)$$

by the induction hypothesis and definition of $[\![\ ]\!]_\rho$, and hence

$$[\![\Gamma]\!]_\rho = ([\![\Gamma]\!]_\rho \sqcap [\![\alpha]\!]_\rho) \sqsubseteq [\![\beta]\!]_\rho$$

by right adjointness of $\Rightarrow$.

If $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \beta$ by ∧-introduction, then

$$[\![\Gamma]\!]_\rho \sqsubseteq [\![\alpha]\!]_\rho \qquad \text{and} \qquad [\![\Gamma]\!]_\rho \sqsubseteq [\![\beta]\!]_\rho$$

by the induction hypothesis, and hence

$$[\![\Gamma]\!]_\rho \sqsubseteq ([\![\alpha]\!]_\rho \sqcap [\![\beta]\!]_\rho) = [\![\alpha \wedge \beta]\!]_\rho$$

because $\sqcap$ is the greatest lower bound w.r.t. $\sqsubseteq$, and by definition of $[\![\ ]\!]_\rho$.

If $\Gamma \vdash_{\mathcal{N}} \alpha$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ by ∧-elimination, then

$$[\![\Gamma]\!]_\rho \sqsubseteq [\![\alpha \wedge \beta]\!]_\rho = [\![\alpha]\!]_\rho \sqcap [\![\beta]\!]_\rho$$

by the induction hypothesis and definition of $[\![\ ]\!]_\rho$, and hence

$$[\![\Gamma]\!]_\rho \sqsubseteq [\![\alpha]\!]_\rho$$

because $\sqcap$ is a lower bound w.r.t. $\sqsubseteq$. Similarly for β. □

We now go back to the motivation that led to the very notion of a Heyting $\sqcap$-algebra, and consider the so-called **Lindenbaum algebra** of Γ, i.e. the set of equivalence classes of formulas under the equivalence relation induced by provable equivalence from Γ.

Theorem 5.2.3 Algebraic Completeness (Jaskowski [1936], Stone [1937], Tarski [1938]) *For any Γ and α,*

$$\Gamma \models_a \alpha \implies \Gamma \vdash_{\mathcal{N}} \alpha.$$

Proof. We show that if $\Gamma \vdash_{\mathcal{N}} \alpha$ fails, then there is a Heyting $\sqcap$-algebra $\mathcal{A}$ and an environment ρ on it such that $[\![\gamma]\!]_\rho = 1$ for all $\gamma \in \Gamma$, but $[\![\alpha]\!]_\rho \neq 1$. This proves the contrapositive of the stated result.

We consider the structure

$$\mathcal{A}_\Gamma = \langle A_\Gamma, \sqsubseteq, =, \sqcap, \Rightarrow, 1 \rangle$$

in which:

1. A_Γ is the set of equivalence classes
$$[\![\beta]\!] = \{\gamma : \Gamma \vdash_{\mathcal{N}} \beta \leftrightarrow \gamma\}$$

2. $\sqsubseteq$ is induced by $\vdash_{\mathcal{N}}$ relatively to Γ, i.e.
$$[\![\beta]\!] \sqsubseteq [\![\gamma]\!] \iff \Gamma, \beta \vdash_{\mathcal{N}} \gamma$$

3. $=$ is induced by provable equivalence relatively to Γ, i.e.
$$[\![\beta]\!] = [\![\gamma]\!] \iff \Gamma \vdash_{\mathcal{N}} (\beta \leftrightarrow \gamma)$$

4. $\sqcap$ is induced by $\wedge$, i.e.
$$[\![\beta]\!] \sqcap [\![\gamma]\!] = [\![\beta \wedge \gamma]\!]$$

5. $\Rightarrow$ is induced by $\rightarrow$, i.e.
$$[\![\beta]\!] \Rightarrow [\![\gamma]\!] = [\![\beta \rightarrow \gamma]\!]$$

6. 1 is the equivalence class of the formulas provable from Γ, i.e.
$$1 = \{\beta : \Gamma \vdash_{\mathcal{N}} \beta\}.$$

Then $\mathcal{A}$ is a Heyting $\sqcap$-algebra, by the discussion motivating the definition of the latter. Since
$$\Gamma \vdash_{\mathcal{N}} \beta \iff [\![\beta]\!] = 1.$$
is the crucial property, we recall the essence of its proof. Basically, the left to right direction comes from the fact that if a formula is provable from Γ, then it is also provable from Γ and any other formula. The right to left direction comes from the fact that if a formula is provably equivalent to a provable formula, then it is itself provable.

In particular, $[\![\gamma]\!] = 1$ for any $\gamma \in \Gamma$, but $[\![\alpha]\!] \neq 1$ because $\Gamma \vdash_{\mathcal{N}} \alpha$ fails by hypothesis. The result then follows by noticing that $[\![\ \ ]\!]$ is actually the canonical algebraic interpretation of formulas on $\mathcal{A}$, w.r.t. the environment defined as
$$\rho(p) = [\![p]\!].$$
We thus have the needed Heyting $\sqcap$-algebra and environment. □

The Algebraic Completeness Theorem provides us with a canonical Heyting $\sqcap$-algebra $\mathcal{A}_\emptyset$, consisting of the equivalence classes of formulas under the equivalence relation induced by intuitionistic provable equivalence.

The Algebraic Soundness Theorem shows that any function from the propositional letters to a Heyting $\sqcap$-algebra $\mathcal{A}$, i.e. any environment on $\mathcal{A}$, can be extended to a homomorphism of Heyting $\sqcap$-algebras from $\mathcal{A}_\emptyset$ to $\mathcal{A}$, i.e. to the canonical algebraic interpretation associated to the environment. This property is concisely expressed by saying that $\mathcal{A}_\emptyset$ is the *free Heyting $\sqcap$-algebra on countably many generators*: namely, the equivalence classes of propositional letters, which are countably many because distinct letters cannot be provably equivalent.

Finite Heyting $\sqcap$-algebras

The next result is the analogue of the Finite Model Property 2.2.11, and by 5.3.10 it is actually a corollary of it (see 18.7.1 for an algebraic proof of it).

Proposition 5.2.4 Finite Model Property (Jaskowski [1936], McKinsey and Tarski [1946]) *If $\Gamma \models_a \alpha$ fails, then there is a finite Heyting algebra $\mathcal{A}$ and an environment ρ on it such that all formulas of Γ are evaluated to* 1 *under it, but α is not.*

Gödel [1933] has proved that no *single* finite Heyting $\sqcap$-algebra is enough, while Jaskowski [1936] has proved that a canonical *countable family* of finite Heyting $\sqcap$-algebras is enough. We prove the first result in 18.7.2. The second result follows from 5.3.10, by considering the finite Heyting $\sqcap$-algebras associated with the finite Kripke models of the canonical family considered in 18.7.1.

5.3 Examples ⋆

We know that, given a lowersemilattice with greatest element, there is *at most one* operation that would satisfy the definition of right adjoint of the g.l.b. operation, but of course there might be *none*. We now look at specific lowersemilattices in which $\sqcap$ does admit a right adjoint, thus providing examples of Heyting $\sqcap$-algebras.

Linear orderings with a greatest element

In a linear ordering the g.l.b. $\sqcap$ of two elements w.r.t. $\sqsubseteq$ is simply the smallest of them, and the adjointness condition becomes

$$(x \sqcap a) \sqsubseteq b \iff x \sqsubseteq (a \Rightarrow b).$$

Proposition 5.3.1 (Gödel [1932]) *Any linear ordering with a greatest element is a Heyting $\sqcap$-algebra.*

Proof. If $a \sqsubseteq b$, then the left-hand-side of the adjointness condition holds for every x, and hence so must the right-hand-side, i.e. $a \Rightarrow b$ is the greatest element 1.

If $b \sqsubset a$, then $(x \sqcap a) \sqsubseteq b$ holds if and only if $x \sqsubseteq b$, and thus

$$x \sqsubseteq b \iff x \sqsubseteq (a \Rightarrow b),$$

i.e. $a \Rightarrow b$ is b.

It is thus enough to let

$$(a \Rightarrow b) = \begin{cases} 1 & \text{if } a \sqsubseteq b \\ b & \text{if } b \sqsubset a \end{cases}$$

to get the right adjoint of $\sqcap$. □

Linear orderings are not enough for completeness because they all satisfy the formula

$$(\alpha \to \beta) \vee (\beta \to \alpha),$$

which is not intuitionistically provable. The next exercises elaborate on this.

Exercises 5.3.2 Linear Heyting $\sqcap$-algebras and Kripke models (Dummett [1959], Horn [1962]) We call a Heyting $\sqcap$-algebra **linear** if its underlying ordering is linear, and similarly for a Kripke model.

a) *Every linear Heyting $\sqcap$-algebra induces an equivalent linear Kripke model.* (Hint: given the linear Heyting $\sqcap$-algebra $\mathcal{L} = \langle L, \sqsubseteq, 1\rangle$ and an environment ρ on it, consider the Kripke model $\mathcal{A} = \langle L, \sqsupseteq, \{\mathcal{A}_x\}_{x \in L}\rangle$, where $\mathcal{A}_x = \{p : \rho(p) \sqsupseteq x\}$. Then

$$x \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; [\![\alpha]\!]_\rho \sqsupseteq x$$

by induction on α, and hence

$$1 \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; [\![\alpha]\!]_\rho = 1.$$

Notice that we inverted the order precisely because we wanted 1 to correspond to truth. And that we could not simply let $\mathcal{A}_x = \{p : \rho(p) = x\}$, because forcing is monotone while ρ needs not be.)

b) *Every linear Kripke model induces an equivalent linear Heyting $\sqcap$-algebra.* (Hint: given the linear Kripke model $\mathcal{A} = \langle P, \sqsubseteq, \{\mathcal{A}_\sigma\}_{\sigma \in P}\rangle$, consider the set of all upward closed subsets of P ordered by inclusion, and the environment

$$\rho(p) = \{\sigma : \sigma \Vdash_{\mathcal{A}} p\}.$$

Then

$$[\![\alpha]\!]_\rho = \{\sigma : \sigma \Vdash_{\mathcal{A}} \alpha\}$$

by induction on α.)

c) *A Heyting $\sqcap$-algebra is linear if and only if it satisfies $(a \Rightarrow b) = 1$ or $(b \Rightarrow a) = 1$, for any a and b.* (Hint: since $1 \sqcap a = a$ and $1 \sqcap b = b$, by adjointness the condition is equivalent to $a \sqsubseteq b$ or $b \sqsubseteq a$, i.e. to linearity.)

d) *A Kripke model is linear if and only if it forces $\alpha \to \beta$ or $\beta \to \alpha$, for any α and β.* (Hint: suppose $\sigma \not\Vdash \alpha \to \beta$ and $\tau \not\Vdash \beta \to \alpha$. Then $\sigma \Vdash \alpha$ but $\sigma \not\Vdash \beta$, and $\tau \Vdash \beta$ but

$\tau \not\Vdash \alpha$. By monotonicity of forcing, neither of σ and τ can be above the other, and the model is not linear.

Conversely, suppose $p \in \mathcal{A}_\sigma - \mathcal{A}_\tau$ and $q \in \mathcal{A}_\tau - \mathcal{A}_\sigma$. Since $p \in \mathcal{A}_\sigma$ but $q \notin \mathcal{A}_\sigma$, $\sigma \not\Vdash p \to q$. Since $q \in \mathcal{A}_\tau$ but $p \notin \mathcal{A}_\tau$, $\tau \not\Vdash q \to p$. Then the model does not force $p \to q$ nor $q \to p$.)

Power sets

In a power set $\mathcal{P}(A)$, consisting of all subsets of a set A ordered under inclusion, the g.l.b. operation is set-theoretical intersection, and the adjointness condition becomes

$$(x \cap a) \subseteq b \iff x \subseteq (a \Rightarrow b).$$

Proposition 5.3.3 *Any power set, ordered under set-theoretical inclusion, is a Heyting $\sqcap$-algebra.*

Proof. Suppose $(x \cap a) \subseteq b$. If $z \in x$, then either $z \notin a$ or $z \in x \cap a$, and hence $z \in b$. Thus $z \in \overline{a} \cup b$, where $\overline{a}$ is the complement of a (relative to A), and $x \subseteq \overline{a} \cup b$. This suggests that

$$(x \cap a) \subseteq b \iff x \subseteq (\overline{a} \cup b),$$

since it proves the left to right direction. Conversely, let $x \subseteq (\overline{a} \cup b)$ and $z \in x \cap a$. Then $z \in b$, and $(x \cap a) \subseteq b$.

It is thus enough to let

$$(a \Rightarrow b) = \overline{a} \cup b.$$

to get the right adjoint of $\cap$. □

Power sets are not enough for completeness since they all satisfy the Law of the Excluded Middle (or its implicational version, Peirce's Law), which is not intuitionistically provable.

Open sets

The example of the power set of a given set A can be generalized by considering not all subsets of A, but only the *open sets* in a topology on A, again ordered under inclusion. Recall that a **topology** $\mathcal{T}$ on A is a collection of subsets of A, called **open sets**, with the following properties:

- $\emptyset$ and A are in $\mathcal{T}$
- $\mathcal{T}$ is closed under arbitrary unions
- $\mathcal{T}$ is closed under finite intersections.

As a typical example of topology, we can take the subsets of the plane that are unions of 'open disks', i.e. circles without their borders.

In general the complement of an open set in a topology is not open, and thus we cannot define the right adjoint of $\cap$ as for power sets. But since

$$(x \cap a) \subseteq b \iff x \subseteq (\overline{a} \cup b)$$

continues to hold, we guess that the following would do:

$$\begin{aligned}(a \Rightarrow b) &= \text{the largest open set contained in } \overline{a} \cup b \\ &= \bigcup \{x : x \in \mathcal{T} \ \wedge \ (x \cap a) \subseteq b\}.\end{aligned}$$

In technical terms, the right-hand-side is called the **interior** of $\overline{a} \cup b$, and is indicated by $(\overline{a} \cup b)^\circ$.

Proposition 5.3.4 (Stone [1937], Tarski [1938]) *The open sets of a topology, ordered under set-theoretical inclusion, form a Heyting $\sqcap$-algebra (called a* **topological** *Heyting algebra).*

Proof. Since the open sets of a topology $\mathcal{T}$ on A ordered by inclusion obviously form a lowersemilattice, with set-theoretical intersection as g.l.b. and A as greatest element, it remains to verify that

$$(x \cap a) \subseteq b \iff x \subseteq (\overline{a} \cup b)^\circ$$

for any open sets x, a and b in $\mathcal{T}$.

The left to right direction holds by definition of interior. For the right to left direction, let $x \subseteq (\overline{a} \cup b)^\circ$ and $z \in x \cap a$. Since $z \in x$, $z \in (\overline{a} \cup b)^\circ$, and hence there is $x_z \in \mathcal{T}$ such that $(x_z \cap a) \subseteq b$ and $z \in x_z$. Since $z \in a$, $z \in x_z \cap a$, and hence $z \in b$. Thus $(x \cap a) \subseteq b$. □

We can thus interpret intuitionistic logic by using open sets. Historically, this is how the first complete interpretation was discovered. The idea underlying it was to *express incompleteness of knowledge by taking as truth-values not points, but whole neighborhoods*.

Tarski [1938] showed that there are *single* topological spaces which are enough for completeness, for example the real line or any n-dimensional euclidean space for $n \geq 1$, with the usual topology (see 18.7.4).

Complete lattices with continuous g.l.b.

We used very little of the properties of open sets in the previous proof, and the following result provides an abstract version.

Theorem 5.3.5 Adjoint Existence (Freyd [1964]) *Any lowersemilattice*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, 1\rangle$$

satisfying the following two conditions is a Heyting $\sqcap$-algebra:

1. *every subset X of A has a l.u.b. $\bigsqcup X$*
2. *$\sqcap$ preserves $\bigsqcup$, i.e. the following $\sqcap\bigsqcup$-distributive law holds:*

$$(\bigsqcup X) \sqcap a = \bigsqcup \{x \sqcap a : x \in X\}$$

Proof. As above, we define

$$(a \Rightarrow b) = \bigsqcup \{x : (x \sqcap a) \sqsubseteq b\},$$

and proves that

$$(x \sqcap a) \sqsubseteq b \iff x \sqsubseteq (a \Rightarrow b).$$

The left to right direction holds by definition of $\Rightarrow$. For the right to left direction, let $x \sqsubseteq (a \Rightarrow b)$. Then

$$\begin{array}{rcl} x \sqcap a & \sqsubseteq & (a \Rightarrow b) \sqcap a \\ & = & (\bigsqcup \{z : (z \sqcap a) \sqsubseteq b\}) \sqcap a \\ & = & \bigsqcup \{z \sqcap a : (z \sqcap a) \sqsubseteq b\} \\ & \sqsubseteq & b, \end{array}$$

respectively because $\sqcap$ preserves the order, by definition of $\Rightarrow$, because $\sqcap$ preserves l.u.b.'s, and because the l.u.b. of a set of elements all less than b is itself less than b. □

Notice that the $\sqsupseteq$ part of condition 2 follows from condition 1, and thus only the $\sqsubseteq$ part needs to be assumed.

The various properties used in the result are well-known. The existence of arbitrary l.u.b.'s makes $\mathcal{A}$ a **complete lattice**, in the sense that g.l.b.'s and l.u.b.'s exist for all subsets: indeed, the g.l.b. of a subset X of A is

$$\bigsqcup \{a : (\forall x \in X)(a \sqsubseteq x)\}.$$

The property of preserving arbitrary l.u.b.'s is *continuity* in the sense of 6.3.7, relative to all l.u.b.'s of arbitrary sets, not only of chains.

The following result is the converse of the previous one.

Proposition 5.3.6 Continuity of G.l.b.'s with Right Adjoints. *In a lowersemilattice*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, 1 \rangle,$$

if $\sqcap$ has a right adjoint $\Rightarrow$, then $\sqcap$ preserves all existing l.u.b.'s.

Proof. We prove that

$$(\bigsqcup X) \sqcap a = \bigsqcup \{x \sqcap a : x \in X\},$$

whenever $\bigsqcup X$ exists.

If $x \in X$, then

$$x \sqsubseteq \bigsqcup X$$

because $\bigsqcup$ is an upper bound, and

$$(x \sqcap a) \sqsubseteq (\bigsqcup X) \sqcap a$$

because $\sqcap$ preserves the order. Thus

$$\bigsqcup \{x \sqcap a : x \in X\} \sqsubseteq (\bigsqcup X) \sqcap a,$$

i.e. $\bigsqcup X \sqcap a$ is an upper bound.

Let now b be any other upper bound to $\bigsqcup \{x \sqcap a : x \in X\}$. For any $x \in X$,

$$(x \sqcap a) \sqsubseteq b,$$

so

$$x \sqsubseteq (a \Rightarrow b)$$

by adjointness,

$$\bigsqcup X \sqsubseteq (a \Rightarrow b)$$

because $\bigsqcup X$ is the least upper bound of X, and

$$(\bigsqcup X \sqcap a) \sqsubseteq b$$

by adjointness, i.e. $\bigsqcup X \sqcap a$ is the least upper bound. □

5.3.5 and 5.3.6 together show that *a complete lattice is a Heyting algebra if and only if it satisfies the $\sqcap\bigsqcup$-distributive law*. Thus the classes of complete Heyting algebras and of complete lattices satisfying the $\sqcap\bigsqcup$-distributive law coincide. However, the two descriptions of the same class stress different aspects: the existence of $\Rightarrow$ in the first case, and the interplay between $\sqcap$ and $\bigsqcup$ in the second. This gives rise to two different 'categories' (a concept that will be introduced in the next chapter), in which the underlying objects are the same but the morphisms preserve, on top of the lattice structure (i.e. $\sqcap$ and $\sqcup$), also the relevant additional aspects. These two categories are called:

- **complete Heyting algebras**, when the morphisms preserve $\Rightarrow$
- **frames**, when the morphisms preserve $\bigsqcup$.

The need to keep the two categories distinct comes from the fact that, although $\Rightarrow$ can be defined in terms of $\sqcap$ and $\bigsqcup$, a morphism preserving the latter does not necessarily preserve the former. What happens here is that, since

$$a \Rightarrow b \ = \ \bigsqcup\{x : x \sqcap a \sqsubseteq b\},$$

if f preserves $\sqcap$ and $\bigsqcup$, then

$$f(a \Rightarrow b) \ = \ \bigsqcup\{f(x) : x \sqcap a \sqsubseteq b\},$$

but

$$f(a) \Rightarrow f(b) \ = \ \bigsqcup\{z : z \sqcap f(a) \sqsubseteq f(b)\}.$$

The relation

$$f(a \Rightarrow b) \sqsubseteq f(a) \Rightarrow f(b)$$

holds automatically, because if $x \sqcap a \sqsubseteq b$, then $f(x) \sqcap f(a) \sqsubseteq f(b))$. But the converse relation is not necessarily true.

Exercise 5.3.7 *There is a frame morphism which is not a morphism of complete Heyting algebras*. (Hint: if $\mathcal{A}$ is a finite linear ordering, a monotone function $f : A \to A$ preserves $\sqcap$ and $\sqcup$, and if it preserves 0 and 1 it is thus a frame morphism. If, for some $b \sqsubset a$, f collapses the closed interval $[b, a]$ into a single element $\neq 1$, then

$$f(a \Rightarrow b) = f(b) \neq 1 = (f(a) \Rightarrow f(b)),$$

and f does not preserve $\Rightarrow$.)

Relationships with Beth-Kripke models $\star$

Partially ordered sets can be given a natural topology, as follows.

Definition 5.3.8 *The* **order topology** $\mathcal{O}_P$ *on a partially ordered set* $(P, \sqsubseteq)$ *consists of the subsets of* P *upward closed w.r.t. to* $\sqsubseteq$. *I.e.* $B \subseteq P$ *is* **open** *if and only if, for every* x *and* y *in* P*:*

$$x \in B \ \wedge \ x \sqsubseteq y \ \Longrightarrow \ y \in B.$$

It is clear that $\mathcal{O}_P$ is a topology, since $\emptyset$ and P are trivially upward closed, and *arbitrary* unions *and* intersections of upward closed sets are still upward closed.

Exercise 5.3.9 A topology $\mathcal{T}$ on X is called $\mathbf{T}_0$ if, for every x and y in X,

$$x \neq y \implies \mathcal{T}_x \neq \mathcal{T}_y,$$

where $\mathcal{T}_x$ is the set of open sets to which x belongs. In other words, if $x \neq y$, then there is an open set containing one of x and y but not the other.

A topology is the order topology associated with some partial ordering if and only if it is T_0 and closed under arbitrary intersections. (Alexandrov [1937]) (Hint: the conditions are sufficient. E.g. if $x \neq y$, then $x \not\sqsubseteq y$ or $y \not\sqsubseteq x$ because $\sqsubseteq$ is antisymmetric, and so y does not belong to the upward closure of x or x does not belong to the upward closure of y. Conversely, given a topological space $\mathcal{T}$, define

$$x \sqsubseteq y \iff \mathcal{T}_x \subseteq \mathcal{T}_y.$$

If $\mathcal{T}$ is T_0, then $\sqsubseteq$ is a partial ordering and it defines an order topology $\mathcal{O}$. An open set A in $\mathcal{T}$ is also open in $\mathcal{O}$. If $\mathcal{T}$ is closed under arbitrary intersections, the opposite holds too, since it does for the upward closure of elements (and these open sets generate the topology). Indeed, given x, its upward closure is

$$\begin{aligned}\{y : x \sqsubseteq y\} &= \{y : \mathcal{T}_x \subseteq \mathcal{T}_y\} \\ &= \{y : (\forall B \in \mathcal{T})(x \in B \Rightarrow y \in B)\} \\ &= \bigcap \{B : B \in \mathcal{T} \wedge x \in B\}.)\end{aligned}$$

Having associated topologies with partial orderings allows us to see any Beth-Kripke model

$$\mathcal{A} = \langle P, \sqsubseteq, \{\mathcal{A}_\sigma\}_{\sigma \in P}\rangle$$

as a model of $\mathcal{N}$ in two different ways:

1. *as a Beth-Kripke model*
   By definition of logical consequence 2.2.4 and the Intuitionistic Soundness Theorem 2.2.5, we know that if $\Gamma \vdash_{\mathcal{N}} \alpha$, then any $\sigma \in P$ forces α whenever it forces all formulas in Γ.

2. *as a topological Heyting $\sqcap$-algebra*
   Here we consider the order topology associated with $(P, \sqsubseteq)$, which is a Heyting $\sqcap$-algebra by 5.3.4. In this case the definition of a canonical algebraic interpretation amounts to:

   $$[\![\alpha]\!]_\rho = \begin{cases} \rho(p) & \text{if } \alpha = p \\ [\![\beta]\!]_\rho \cap [\![\gamma]\!]_\rho & \text{if } \alpha = \beta \wedge \gamma \\ (\overline{[\![\beta]\!]_\rho} \cup [\![\gamma]\!]_\rho)^\circ & \text{if } \alpha = \beta \to \gamma. \end{cases}$$

   By definition of algebraic consequence 5.2.1 and the Algebraic Soundness Theorem 5.2.2, we know that if $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$, then

   $$([\![\gamma_1]\!]_\rho \cap \cdots \cap [\![\gamma_n]\!]_\rho) \subseteq [\![\alpha]\!]_\rho$$

   for every environment ρ.

We now show that the former notion is actually a special case of the latter, for a particular environment ρ.

Proposition 5.3.10 Canonical Topological Interpretation Induced by Forcing. *Given a Beth-Kripke model*

$$\mathcal{A} = \langle P, \sqsubseteq, \{\mathcal{A}_\sigma\}_{\sigma \in P} \rangle,$$

the environment

$$\rho(p) = \{\sigma : \sigma \Vdash_{\mathcal{A}} p\}$$

produces on the topological Heyting $\sqcap$*-algebra*

$$\langle \mathcal{O}_P, \subseteq, \cap, \Rightarrow, P \rangle$$

a canonical algebraic interpretation that coincides with forcing, in the sense that for every formula α

$$\llbracket \alpha \rrbracket_\rho = \{\sigma : \sigma \Vdash_{\mathcal{A}} \alpha\}.$$

Proof. We prove that

$$\sigma \in \llbracket \alpha \rrbracket_\rho \iff \sigma \Vdash_{\mathcal{A}} \alpha$$

by induction on α.

If $\alpha = p$, then this holds by definition of ρ. Notice that ρ really is an environment, since the set of states forcing a given formula, p in particular, is upward closed by monotonicity of forcing.

If $\alpha = \beta \wedge \gamma$, then

$$\begin{aligned} \sigma \in \llbracket \beta \wedge \gamma \rrbracket_\rho &\iff \sigma \in (\llbracket \beta \rrbracket_\rho \cap \llbracket \gamma \rrbracket_\rho) \\ &\iff \sigma \in \llbracket \beta \rrbracket_\rho \ \wedge\ \sigma \in \llbracket \gamma \rrbracket_\rho \\ &\iff \sigma \Vdash_{\mathcal{A}} \beta \ \wedge\ \sigma \Vdash_{\mathcal{A}} \gamma \\ &\iff \sigma \Vdash_{\mathcal{A}} \beta \wedge \gamma \end{aligned}$$

by definition of $\llbracket\ \rrbracket$, induction hypothesis and definition of forcing for $\wedge$.

If $\alpha = \beta \to \gamma$, then $\llbracket \beta \to \gamma \rrbracket_\rho$ is the interior of $\overline{\llbracket \beta \rrbracket_\rho} \cup \llbracket \gamma \rrbracket_\rho$, i.e. the union of all open sets contained in it, and

$$\sigma \Vdash_{\mathcal{A}} \beta \to \gamma \iff (\forall \tau \sqsupseteq \sigma)(\tau \Vdash_{\mathcal{A}} \beta \ \Rightarrow\ \tau \Vdash_{\mathcal{A}} \gamma).$$

If $\sigma \in \llbracket \beta \to \gamma \rrbracket_\rho$, then σ belongs to an open set contained in $\overline{\llbracket \beta \rrbracket_\rho} \cup \llbracket \gamma \rrbracket_\rho$ by definition of interior. In particular, all extensions τ of σ are in the same set, because open sets are upward closed. Then $\tau \in \overline{\llbracket \beta \rrbracket_\rho} \cup \llbracket \gamma \rrbracket_\rho$, and hence

$$\tau \in \llbracket \beta \rrbracket_\rho \implies \tau \in \llbracket \gamma \rrbracket_\rho.$$

By the induction hypothesis,

$$\tau \Vdash_{\mathcal{A}} \beta \implies \tau \Vdash_{\mathcal{A}} \gamma.$$

Since τ is an arbitrary extension of σ, $\sigma \Vdash_{\mathcal{A}} \beta \to \gamma$ by definition of forcing.

Conversely, if $\sigma \Vdash_{\mathcal{A}} \beta \to \gamma$, then

$$\tau \Vdash_{\mathcal{A}} \beta \implies \tau \Vdash_{\mathcal{A}} \gamma$$

for all extensions τ of σ. By the induction hypothesis,

$$\tau \in \llbracket \beta \rrbracket_\rho \implies \tau \in \llbracket \gamma \rrbracket_\rho,$$

i.e. all extensions of σ are in $\overline{\llbracket \beta \rrbracket_\rho} \cup \llbracket \gamma \rrbracket_\rho$. But the set of all extensions of σ is upward closed, and hence an open set. Then σ belongs to an open set contained in $\overline{\llbracket \beta \rrbracket_\rho} \cup \llbracket \gamma \rrbracket_\rho$, and hence is in its interior $\llbracket \beta \to \gamma \rrbracket_\rho$. □

Now the fact that $\mathcal{A}$ is a Beth-Kripke model, i.e. that, for $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$,

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies (\forall \sigma \in P)(\sigma \Vdash_{\mathcal{A}} \gamma_1 \wedge \cdots \wedge \gamma_n \implies \sigma \Vdash_{\mathcal{A}} \alpha),$$

can be translated as

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies (\llbracket \gamma_1 \rrbracket_\rho \cap \cdots \cap \llbracket \gamma_n \rrbracket_\rho) \subseteq \llbracket \alpha \rrbracket_\rho,$$

and it is a special case of the fact that the order topology associated with $\mathcal{A}$ is a Heyting $\sqcap$-algebra.

In particular, *the Algebraic Soundness Theorem is a strengthening of the Intuitionistic Soundness Theorem*, since it exhibits a larger class of interpretations for provable formulas, and *the Intuitionistic Completeness Theorem is a strengthening of the Algebraic Completeness Theorem*, since it exhibits a smaller class of counterexamples for disprovable formulas.

Both strengthenings are real by 5.3.9, which shows that not every topology arises from a partial ordering, and hence that not every topological Heyting algebra arises from a Beth-Kripke model.

Exercises 5.3.11 Forcing induced by a canonical interpretation.

a) *Given a Heyting $\sqcap$-algebra*

$$\mathcal{A} = \langle A, \sqsubseteq, \sqcap, \Rightarrow, 1 \rangle,$$

any environment ρ on it produces a Beth-Kripke model

$$\mathcal{B} = \langle P, \sqsubseteq, \{\mathcal{B}_\sigma\}_{\sigma \in P} \rangle$$

whose forcing coincides with the canonical algebraic interpretation, in the sense that, for every formula α,

$$\sigma \Vdash \alpha \iff [\![\alpha]\!]_\rho \in \sigma.$$

(Hint: let $P = \mathcal{F}_A$ be the set of filters on A, $\sqsubseteq$ be the set-theoretical inclusion, $\mathcal{B}_\sigma$ be $\{p : [\![p]\!]_\rho \in \sigma\}$, and proceed by induction on α.

To treat $\rightarrow$, we need to show that

$$(\forall \tau \supseteq \sigma)([\![\beta]\!] \in \tau \Longrightarrow [\![\gamma]\!] \in \tau) \iff ([\![\beta]\!] \Rightarrow [\![\gamma]\!]) \in \sigma.$$

The right to left direction follows from 5.1.11. For the left to right direction, given $[\![\sigma]\!]$, consider the filter τ generated by $\sigma \cup \{[\![\beta]\!]\}$. By hypothesis $[\![\gamma]\!] \in \tau$, i.e. there is $a \in \sigma$ such that $a \sqcap [\![\beta]\!] \sqsubseteq [\![\gamma]\!]$, and $a \sqsubseteq ([\![\beta]\!] \Rightarrow [\![\gamma]\!])$ by adjointness. By upward closure of σ, $([\![\beta]\!] \Rightarrow [\![\gamma]\!]) \in \sigma$.)

b) *The Beth-Kripke model* $\mathcal{F}_{\mathcal{O}_P}$ *associated with a Heyting* $\sqcap$*-algebra* $\mathcal{O}_P$ *associated with a Beth-Kripke model* P *is isomorphic to* P. (Hint: every filter on $\mathcal{O}_P$ is principal, because $\mathcal{O}_P$ is closed under arbitrary intersections.)

c) *The Heyting* $\sqcap$*-algebra* $\mathcal{O}_{\mathcal{F}_A}$ *associated with a Beth-Kripke model* $\mathcal{F}_A$ *associated with a Heyting* $\sqcap$*-algebra* A *is not in general isomorphic to* A. (Hint: by 18.3.4, not every Heyting $\sqcap$-algebra is isomorphic to a topological one.)

5.4 Representation Theorems ⋆

We now ask how far the examples of Heyting $\sqcap$-algebras produced in Section 5.3 are typical. The results of this section are generalizations of the ones of Section 20.3 on Boolean algebras. We thus suggest the reader to read that section first, to see similar proofs in a simpler setting.

Lindenbaum algebras

Theorem 5.4.1 First Representation for Heyting $\sqcap$-algebras (Tarski [1935a]) *Any Heyting* $\sqcap$*-algebra is isomorphic to a Lindenbaum algebra for the Implicational Calculus with Conjunction.*

Proof. Given a Heyting $\sqcap$-algebra $\mathcal{A}$, consider a language for the Implicational Calculus with Conjunction with as many propositional letters as there are elements of A, and define an environment

$$\rho : \text{Propositional Letters} \longrightarrow A$$

which is onto A. Then let $[\![\alpha]\!]_\rho$ be the canonical algebraic interpretation of α w.r.t. ρ.

By definition of canonical algebraic interpretation and the Algebraic Soundness Theorem, $[\![\ ]\!]_\rho$ induces a homomorphism of Heyting $\sqcap$-algebras between the

Lindenbaum algebra of the given propositional language and $\mathcal{A}$. Moreover, the homomorphism is onto because A is already covered by ρ.

To get an isomorphism, we simply identify elements in the Lindenbaum algebra that get mapped to the same element of A. This can be done canonically, by defining a set of formulas Γ as follows:

$$\alpha \in \Gamma \iff [\![\alpha]\!]_\rho = 1.$$

By definition of canonical algebraic interpretation and the Algebraic Soundness Theorem, Γ is deductively closed, i.e.

$$\Gamma \vdash_{\mathcal{N}} \alpha \iff [\![\alpha]\!]_\rho = 1,$$

and

$$\alpha \leftrightarrow \beta \in \Gamma \iff [\![\alpha]\!]_\rho = [\![\beta]\!]_\rho.$$

Then the Lindenbaum algebra associated with Γ is actually isomorphic to $\mathcal{A}$. □

Open sets

The Algebraic Soundness Theorem shows that Heyting $\sqcap$-algebras provide us with a general algebraic formulation of the notion of model for the Implicational Calculus with Conjunction. Conversely, we now show that the proof of the Intuitionistic Completeness Theorem contains in a nutshell the ideas needed for a general representation theorem for Heyting $\sqcap$-algebras which shows that, in a precise sense, *the only Heyting $\sqcap$-algebras are the topological ones and their subalgebras*.

Theorem 5.4.2 Second Representation for Heyting $\sqcap$-algebras (Stone [1937], McKinsey and Tarski [1946]) *Any Heyting $\sqcap$-algebra is isomorphic to a subalgebra of an algebra of open sets of a topology.*

Proof. In 5.1.10 we introduced the notion of a filter of a Heyting $\sqcap$-algebra, by modelling it on the properties of sets of formulas needed to prove the Intuitionistic Completeness Theorem (2.2.6 and 4.2.2). We now exploit the analogy, by considering the set $\mathcal{F}_A$ of all filters on A, and the function f defined as follows:

$$\begin{aligned} f(x) &= \text{the set of all filters containing } x \\ &= \{F : F \in \mathcal{F}_A \wedge x \in F\}. \end{aligned}$$

Notice that f is a function from A to the power set $\mathcal{P}(\mathcal{F}_A)$ of $\mathcal{F}_A$, which is a lowersemilattice w.r.t. inclusion, with the following properties:

1. *f preserves $\sqsubseteq$*
   Given $x \sqsubseteq y$, let F be a filter containing x. By upward closure, F also contains y. Thus
   $$x \sqsubseteq y \implies f(x) \subseteq f(y).$$

Actually, the converse implication holds as well: if $f(x) \subseteq f(y)$, then the principal filter generated by x is in $f(y)$, i.e. it contains y, and hence $x \sqsubseteq y$.

2. *f preserves $\sqcap$*
   Given x, y and a filter F, if F contains $x \sqcap y$, then it also contains x (by upward closure, since $x \sqcap y \sqsubseteq x$) and y (similarly), and so $f(x \sqcap y) \subseteq f(x) \cap f(y)$.

   Conversely, if F contains x and y, then it also contains $x \sqcap y$ by definition of filter, and so $f(x) \cap f(y) \subseteq f(x \sqcap y)$. Thus

   $$f(x \sqcap y) = f(x) \cap f(y).$$

3. *f preserves 1*
   1 is in any filter F, since it is the greatest element of A and F is upward closed. Thus

   $$f(1) = \mathcal{F}_A.$$

4. *there is a topology on $\mathcal{F}_A$ w.r.t. which f preserves $\Rightarrow$*
   Since right adjointness can be presented equationally in terms of $\sqsubseteq$ and $\sqcap$ (5.1.8), which are preserved by f, $\Rightarrow$ is automatically preserved whenever the definition of right adjointness of $\cap$ in $\mathcal{P}(\mathcal{F}_A)$ only requires consideration of elements in the range of f. This can be obtained by defining the topology on $\mathcal{F}_A$ as the *topology generated by* $f(A)$, i.e. as the smallest topology for which all elements in the range of f are open. More explictly, by defining the open sets as arbitrary unions of elements of $f(A)$. Notice that $f(A)$ is already closed under finite intersections by part 2 above, and it contains $\mathcal{F}_A$ by part 3.

   To formally verify that f preserves $\Rightarrow$, by uniqueness of right adjoints we only need to show that, for any open set x,

   $$x \cap f(a) \subseteq f(b) \iff x \subseteq f(a \Rightarrow b).$$

   By definition of topology generated by $f(A)$, there is a subset B of A such that $x = \bigcup_{z \in B} f(z)$. Then

   $$\begin{array}{rl} & x \cap f(a) \subseteq f(b) \\ \iff & (\bigcup\limits_{z \in B} f(z)) \cap f(a) \subseteq f(b) \\ \iff & \bigcup\limits_{z \in B} (f(z) \cap f(a)) \subseteq f(b) \\ \iff & (\forall z \in B)[f(z) \cap f(a) \subseteq f(b)] \\ \iff & (\forall z \in B)[f(z \sqcap a) \subseteq f(b)] \end{array}$$

$$
\begin{aligned}
&\iff (\forall z \in B)[z \sqcap a \sqsubseteq b] \\
&\iff (\forall z \in B)[z \sqsubseteq a \Rightarrow b] \\
&\iff (\forall z \in B)[f(z) \subseteq f(a \Rightarrow b)] \\
&\iff (\bigcup_{z \in B} f(z)) \subseteq f(a \Rightarrow b) \\
&\iff x \subseteq f(a \Rightarrow b),
\end{aligned}
$$

by definition of x, set-theoretical $\cap\bigcup$-distributive law, definition of adjointeness, and preservation properties of f.

5. *f is one-one*
Given x and y, if $x \neq y$, then (since $\sqsubseteq$ is a partial ordering, in particular antisymmetric) $x \not\sqsubseteq y$ or $y \not\sqsubseteq x$. If $x \not\sqsubseteq y$, then the upward closure of x is obviously a filter containing x but not y, and thus $f(x) \neq f(y)$. Similarly when $y \not\sqsubseteq x$. Thus
$$x \neq y \implies f(x) \neq f(y).$$

We have thus proved that f induces an isomorphism of Heyting $\sqcap$-algebras between A and $f(A)$, where $f(A)$ is intended as a subalgebra of the topological algebra $\mathcal{P}(\mathcal{F}_A)$. □

The set $\mathcal{F}_A$ of filters of A and its topology generated by $f(A)$ are respectively called the **Stone space** of A and the **Stone topology** associated with it. Thus the previous result can be reformulated in the following more informative way, usually referred to as the **Stone Representation Theorem**: *any Heyting $\sqcap$-algebra is isomorphic to a subalgebra of the algebra of open sets of its Stone space.*

This formulation of the Stone Representation Theorem raises two complementary questions:

- Which Heyting $\sqcap$-algebras are isomorphic not only to a *subalgebra*, but to a *full algebra* of open sets of a topology?

- Which topologies arise as Stone topologies of Heyting $\sqcap$-algebras?

We will answer the first question in 18.3.4, and a dual version of the second in 18.5.7.

æ

Chapter 6

Cartesian Closed Categories

In the present chapter we pursue the idea that the proof of an implication is a *function* from the set of proofs of (the conjunction of) its premises to the set of proofs of its conclusion. More precisely, we describe a view of logic in which formulas are interpreted as sets (of proofs), proofs as functions between formulas, and deduction rules as operations on functions.

We will provide the needed categorical background, and refer to MacLane [1971], Arbib and Manes [1975], Goldblatt [1979], Lawvere and Schanuel [1991], and Pierce [1991] for more detailed treatment. Similarly, Lambek and Scott [1986], and Asperti and Longo [1991] are the references for detailed treatmentes of the connections with typed lambda calculus.

6.1 Cartesian Closed Categories

Categorical Models

To describe the interplay of $\vdash_{\mathcal{N}}$, $\wedge$ and $\rightarrow$ we consider categorical structures $\mathcal{C}$ with:

- A class[1] $Ob_{\mathcal{C}}$ of sets called **objects**, intended to interpret formulas.
- A binary class function $Hom_{\mathcal{C}}$ that associates to every pair of objects A and B a (possibly empty) set $Hom_{\mathcal{C}}(A, B)$ of functions with domain A and

[1] The word *class* is taken here in a technical sense, and not as a synonym of *set*. Basically, unrestricted predicates define classes, while predicates restricted to sets define sets. Thus sets are defined inductively, and the inductive conditions for the definition are given by the usual axioms of Set Theory. We can also think of sets as ‘small classes’, or classes belonging to some (other) class. Then ‘large’ classes, called ‘proper’, are classes that are not sets. Considering classes of objects as opposed to sets will allow us to deal, in the following, with categories such as **Set**, whose collection of objects constitutes a proper class.

codomain B called **morphisms**, intended to interpret proofs from premises and hence to model $\vdash_{\mathcal{N}}$. For convenience, we will often write $f : A \to B$ for $f \in \mathit{Hom}(A, B)$.

- Two binary class functions $\times_{\mathcal{C}}$ and $\Rightarrow_{\mathcal{C}}$, intended to model $\wedge$ and $\to$, respectively.

Such a structure will be used to define an interpretation function $[\![\ ]\!]^{\mathcal{C}}_{\rho}$ for formulas, relative to a given interpretation (called *environment*) of the propositional letters, i.e. to a function ρ from the set of all propositional letters to the underlying class $\mathit{Ob}_{\mathcal{C}}$ of $\mathcal{C}$.

As usual, we will drop the superscript $\mathcal{C}$ or the subscript ρ when no confusion arises.

Definition 6.1.1 Canonical Categorical Interpretation. *Given a structure*

$$\mathcal{C} = \langle \mathit{Ob}, \times, \Rightarrow \rangle$$

and an **environment** ρ *on it, i.e. a function*

$$\rho : \mathit{Propositional\ Letters} \longrightarrow \mathit{Ob},$$

we define the **canonical categorical interpretation** $[\![\ ]\!]_{\rho}$ *by induction on formulas, as follows:*

$$[\![\alpha]\!]_{\rho} = \begin{cases} \rho(p) & \textit{if } \alpha = p \\ [\![\beta]\!]_{\rho} \times [\![\gamma]\!]_{\rho} & \textit{if } \alpha = \beta \wedge \gamma \\ [\![\beta]\!]_{\rho} \Rightarrow [\![\gamma]\!]_{\rho} & \textit{if } \alpha = \beta \to \gamma. \end{cases}$$

By induction, $[\![\alpha]\!]_{\rho} \in \mathit{Ob}$ for every α. Having the notion of interpretation, we define the notion of model by mirroring $\vdash_{\mathcal{N}}$, modulo the identification of finite sets of premises with their conjunctions. Since $\alpha \vdash_{\mathcal{N}} \beta$ says that there is a proof of β from α, and we interpret proofs as morphisms, then the relation $\vdash_{\mathcal{N}}$ will be modelled by the fact that the appropriate *Hom* set is not empty.

Definition 6.1.2 *A* **categorical model** *of Implicational Calculus with Conjunction is a structure*

$$\mathcal{C} = \langle \mathit{Ob}, \mathit{Hom}, \times, \Rightarrow \rangle$$

such that, for every $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ *and* α,

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies (\forall \rho)(\mathit{Hom}([\![\gamma_1]\!]_{\rho} \times \cdots \times [\![\gamma_n]\!]_{\rho}, [\![\alpha]\!]_{\rho}) \neq \emptyset).$$

Categories

Our next goal is to determine conditions ensuring that $\mathcal{C}$ is a categorical model. The discussion proceeds as in Section 5.1.

Reflexivity and transitivity of $\vdash_{\mathcal{N}}$ show that the *Hom* sets must contain the identity functions, and be closed under composition.

Definition 6.1.3 (Eilenberg and MacLane [1945]) *A structure*

$$\mathcal{C} = \langle Ob, Hom \rangle$$

is a **locally small, concrete category** *if the following hold, for every object A, B and C:*

1. *$id_A \in Hom(A, A)$, where id_A is the identity function on A defined by:*

$$id_A(x) = x;$$

2. *if $f \in Hom(A, B)$ and $g \in Hom(B, C)$, then $g \circ f \in Hom(A, C)$, where $\circ$ is the composition operator defined by:*

$$(f \circ g)(x) = f(g(x)).$$

C is a **small concrete category** *if, in addition, Ob is a set.*

An intuitive way of reformulating the rules just given is in terms of a deductive system, with axioms

$$id_A : A \to A,$$

and a deduction rule

$$\frac{f : A \to B \qquad g : B \to C}{g \circ f : A \to C.}$$

In the previous definition, the word *concrete* refers to the fact that we required each $A \in Ob$ to be a set, and each $f \in Hom(A, B)$ to be a set-theoretical function. The words *locally small* refer instead to the fact that we required each $Hom(A, B)$ to be a set. In the main text we will only deal with locally small, concrete categories, and thus drop the qualifiers. In the exercises we will instead introduce and use a more general notion.

Exercises 6.1.4 Abstract categories. We consider a structure $\mathcal{C} = \langle Ob, Hom \rangle$ as above, but without supposing any of the following: the objects, i.e. the elements of Ob, are sets; the morphisms, i.e. the elements of $Hom(A, B)$, are set-theoretical functions; and the collections of morphisms, i.e. the $Hom(A, B)$ themselves, are sets.

Such a structure is called an **abstract category** if the following hold, for every object A, B and C:

1. there is an operator $\circ$ called *composition*, such that:
   (a) if $f \in Hom(A, B)$ and $g \in Hom(B, C)$, then $g \circ f \in Hom(A, C)$
   (b) $\circ$ is associative, i.e.
   $$f \circ (g \circ h) = (f \circ g) \circ h,$$
2. there is a morphism $id_A \in Hom(A, A)$ called the identity on A, such that for every morphism $f \in Hom(A, B)$
$$f \circ id_A = f \qquad \text{and} \qquad id_B \circ f = f.$$

An abstract category is **locally small** if, for every pair of objects A and B, $Hom(A, B)$ is a set. A locally small category is **small** if Ob is a set.

A concrete category is an abstract category. (Hint: the set-theoretical identities and composition obviously satisfy the required additional properties.)

Antisymmetry of $\vdash_{\mathcal{N}}$ suggests the following notion, which plays the role of equality between objects of a category.

Definition 6.1.5 *In a category, two objects A and B are* **isomorphic** *(written* $\boldsymbol{A \simeq B}$*) if there are morphisms $f : A \to B$ and $g : B \to A$ such that*
$$g \circ f = id_A \qquad \textit{and} \qquad f \circ g = id_B.$$
In that case, f is called an **isomorphism** *of A and B.*

Terminal object

When $\vdash_{\mathcal{N}}$ is interpreted as a partial ordering, the existence of formulas provable without assumptions is translated in the existence of a greatest element 1. When we interpret $\vdash_{\mathcal{N}}$ as the existence of morphisms, we notice that formulas provable without assumptions are provable from *any* assumption, and this suggests the consideration of an object 1 such that $Hom(A, 1) \neq \emptyset$, for any object A. To uniquely determine 1 we need something more: not only existence of a morphism, but uniqueness of it.

Definition 6.1.6 *In a category*
$$\mathcal{C} = \langle Ob, Hom \rangle,$$
an object 1 is called a **terminal object** *if there is exactly one morphism t_A in $Hom(A, 1)$, for any object A.*

In terms of the deductive system introduced above, we have further axioms:
$$t_A : A \to 1$$

Proposition 6.1.7 *A terminal object is unique up to isomorphism, when it exists.*

Proof. If A and B are both terminal, there exist morphisms $f : A \to B$ and $g : B \to A$. Then $g \circ f$ and id_A are both morphisms from A into A, and thus they must be equal because A is terminal, i.e. $g \circ f = id_A$. Similarly, $f \circ g = id_B$. Thus A and B are isomorphic. □

Exercise 6.1.8 Initial object. In a category, an object 0 is called an **initial object** if there is exactly one morphism in $Hom(0, A)$, for any object A.

An initial object is unique up to isomorphism, when it exists.

Exercises 6.1.9 Opposite category. Given an abstract category $\mathcal{C}$, consider the structure $\mathcal{C}^{op}$ such that:

- the objects of $\mathcal{C}^{op}$ are the same as the objects of $\mathcal{C}$
- the morphisms of $\mathcal{C}^{op}$ are the morphisms of $\mathcal{C}$ with inverted domains and codomains, i.e. $Hom_{\mathcal{C}^{op}}(A, B) = Hom_{\mathcal{C}}(B, A)$.

a) *$\mathcal{C}^{op}$ is an abstract category, called the* **opposite** *of $\mathcal{C}$.*

b) *The initial object in $\mathcal{C}$ is the terminal object in $\mathcal{C}^{op}$.* This gives an alternative proof of uniqueness of initial objects.

Products

We now interpret $\vdash_{\mathcal{N}}$ as the existence of morphisms, and $\wedge$ by $\times$.

The $\wedge$-elimination rules show that, whenever there is a morphism to $A \times B$, there must be morphisms to A and B too. The easiest way to ensure this is to take advantage of closure under composition, and simply require the existence of morphisms from $A \times B$ to A and B.

The $\wedge$-introduction rule shows that, whenever there are morphisms to both A and B, there must be a morphism to $A \times B$, too.

As already for 1, purely existential properties would not be enough to determine $A \times B$ up to isomorphim, and thus we also add a uniqueness condition.

Definition 6.1.10 *In a category*

$$\mathcal{C} = \langle Ob, Hom \rangle,$$

an operation $\times$ is called a **product** *if, for every pair of objects A and B:*

1. *there are morphisms $l_{A,B} : A \times B \to A$ and $r_{A,B} : A \times B \to B$ called, respectively, the left and right* **projections**
2. *for every object C and any pair $f : C \to A$ and $g : C \to B$ of morphisms from C to A and B, there is a unique morphism*

$$\langle f, g \rangle_{C,A,B} : C \to A \times B$$

such that

$$f = l_{A,B} \circ \langle f, g\rangle_{C,A,B} \qquad \textit{and} \qquad g = r_{A,B} \circ \langle f, g\rangle_{C,A,B}.$$

This can be represented pictorially by saying that in the following diagram

$$\begin{tikzcd}
& C \arrow[dl, "f"'] \arrow[d, dotted, "\langle f{,}g\rangle"] \arrow[dr, "g"] & \\
A & A \times B \arrow[l, "l"'] \arrow[r, "r"] & B
\end{tikzcd}$$

all morphisms from the same domain to the same codomain are equal (the dotted arrow only focuses the attention on the principal morphism). We say in such cases that *the diagram commutes*, or that it is *commutative.*

In terms of the deductive system introduced above, we have further axioms

$$l_{A,B} : A \times B \to A \qquad \text{and} \qquad r_{A,B} : A \times B \to B,$$

and an additional rule

$$\frac{f : C \to A \quad g : C \to B}{\langle f, g\rangle : C \to A \times B.}$$

While the latter literally corresponds to $\wedge$-introduction, $\wedge$-elimination can be stated as the pair of derived rules:

$$\frac{f : C \to A \times B}{l_{A,B} \circ f : C \to A} \qquad \text{and} \qquad \frac{f : C \to A \times B}{r_{A,B} \circ f : C \to B.}$$

Proposition 6.1.11 *The product of two objects is unique up to isomorphism, when it exists.*

Proof. Suppose $A \times^* B$ also satisfies the properties, i.e.

1. there are morphisms $l^*_{A,B} : A \times^* B \to A$ and $r^*_{A,B} : A \times^* B \to B$
2. for every object C and any pair $f : C \to A$ and $g : C \to B$ of morphisms from C to A and B, there is a unique morphism $\langle f, g\rangle^* : C \to A \times^* B$ such that
$$f = l^*_{A,B} \circ \langle f, g\rangle^* \qquad \text{and} \qquad g = r^*_{A,B} \circ \langle f, g\rangle^*.$$

Then, by letting $C = A \times^* B$, $f = l^*_{A,B}$ and $g = r^*_{A,B}$ in the definition of $A \times B$, and $C = A \times B$, $f = l_{A,B}$ and $g = r_{A,B}$ in the definition of $A \times^* B$, we have that

$$\langle l^*_{A,B}, r^*_{A,B}\rangle \circ \langle l_{A,B}, r_{A,B}\rangle^* = id_{A\times B}$$

and

$$\langle l_{A,B}, r_{A,B}\rangle^* \circ \langle l^*_{A,B}, r^*_{A,B}\rangle = id_{A\times^* B}$$

by the uniqueness condition, and thus $A \times B$ and $A \times^* B$ are isomorphic. □

Exercises 6.1.12 The following hold in any category with terminal object and products.

a) $A \times B \simeq B \times A$.

b) $A \times (B \times C) \simeq (A \times B) \times C$.

c) $A \simeq A \times 1$.

d) $\langle f \circ h, g \circ h\rangle = \langle f, g\rangle \circ h$. (Hint: consider the following diagram.

$$\begin{tikzcd}
 & D \arrow[ddl, "f\circ h"'] \arrow[d, "h"] \arrow[ddr, "g\circ h"] & \\
 & C \arrow[dl, "f"'] \arrow[d, "\langle f,g\rangle"] \arrow[dr, "g"] & \\
A & A \times B \arrow[l] \arrow[r] & B.)
\end{tikzcd}$$

e) $\langle l_{A,B}, r_{A,B}\rangle = id_{A\times B}$.

Exercises 6.1.13 Given an abstract category $\mathcal{C}$ and two objects A and B, consider the structure $\mathcal{C}^{\times}_{A,B}$ such that:

- the objects of $\mathcal{C}^{\times}_{A,B}$ are the triples $\langle C, f, g\rangle$, where C is an object of $\mathcal{C}$, and $f : C \to A$ and $g : C \to B$ are two morphisms of $\mathcal{C}$:

$$\begin{tikzcd}
A & C \arrow[l, "f"'] \arrow[r, "g"] & B
\end{tikzcd}$$

- the morphisms $h : (C_1, f_1, g_1) \to (C_2, f_2, g_2)$ of $\mathcal{C}^{\times}_{A,B}$ are the morphisms $h : C_1 \to C_2$ of $\mathcal{C}$ such that $f_1 = f_2 \circ h$ and $g_1 = g_2 \circ h$:

$$\begin{tikzcd}
 & C_1 \arrow[dl, "f_1"'] \arrow[d, dotted, "h"] \arrow[dr, "g_1"] & \\
A & C_2 \arrow[l, "f_2"'] \arrow[r, "g_2"] & B.
\end{tikzcd}$$

a) *$\mathcal{C}^{\times}_{A,B}$ is an abstract category.*

b) *The product $A \times B$ in $\mathcal{C}$ is the terminal object in $\mathcal{C}^{\times}_{A,B}$.* This gives an alternative proof of uniqueness of products in $\mathcal{C}$.

For convenience, we introduce the following notation.

Definition 6.1.14 *Given $f : A \to C$ and $g : B \to D$, let*

$$f \times g : A \times B \to C \times D$$

be the function defined by:

$$f \times g = \langle f \circ l_{A,B}, g \circ r_{A,B} \rangle.$$

As usual, we can just say that the following diagram commutes:

$$\begin{tikzcd}
A \arrow[d, "f"'] & A \times B \arrow[l, "l_{A,B}"'] \arrow[r, "r_{A,B}"] \arrow[dl, "f \circ l_{A,B}"] \arrow[d, dotted, "f \times g"] \arrow[dr, "g \circ r_{A,B}"'] & B \arrow[d, "g"] \\
C & C \times D \arrow[l, "l_{C,D}"'] \arrow[r, "r_{C,D}"] & D.
\end{tikzcd}$$

In terms of the deductive system introduced above, this corresponds to having a derived rule

$$\frac{f : A \to C \qquad g : B \to D}{f \times g : A \times B \to C \times D.}$$

Exercises 6.1.15 The following hold in any category with products.

a) $id_A \times id_B = id_{A \times B}$.

b) $(f \times h) \circ (g \times k) = (f \circ g) \times (h \circ k)$. (Hint: consider the following diagram.

$$\begin{tikzcd}
A \arrow[d, "g"'] & A \times B \arrow[l] \arrow[r] \arrow[d, "g \times k"] & B \arrow[d, "k"] \\
C \arrow[d, "f"'] & C \times D \arrow[l] \arrow[r] \arrow[d, "f \times h"] & D \arrow[d, "h"] \\
E & E \times F \arrow[l] \arrow[r] & F.)
\end{tikzcd}$$

c) $(f \times h) \circ \langle g, k \rangle = \langle f \circ g, h \circ k \rangle$. (Hint: consider the following diagram.

$$\begin{tikzcd}
 & A \arrow[dl, "g"'] \arrow[d, "\langle g,k \rangle"] \arrow[dr, "k"] & \\
B \arrow[d, "f"'] & B \times C \arrow[l] \arrow[r] \arrow[d, "f \times h"] & C \arrow[d, "h"] \\
D & D \times E \arrow[l] \arrow[r] & E.)
\end{tikzcd}$$

d) $(f \times id) \circ \langle id, k \rangle = \langle f, g \rangle$. (Hint: from part c).)

Exercises 6.1.16 Given an abstract category $\mathcal{C}$, consider the structure $\mathcal{C}^{\rightarrow}$ such that:

- the objects of $\mathcal{C}^{\rightarrow}$ are the morphisms $f : A \rightarrow B$ of $\mathcal{C}$;
- the morphisms $(h_1, h_2) : f \rightarrow g$ of $\mathcal{C}^{\rightarrow}$, where $f : A \rightarrow B$ and $g : C \rightarrow D$, are the pairs of morphisms $h_1 : A \rightarrow C$ and $h_2 : B \rightarrow D$, such that $g \circ h_1 = h_2 \circ f$:

$$\begin{tikzcd}
A \arrow[r, "f"] \arrow[d, "h_1"'] & B \arrow[d, "h_2"] \\
C \arrow[r, "g"] & D.
\end{tikzcd}$$

a) *$\mathcal{C}^{\rightarrow}$ is an abstract category.*

b) *$f \times g$ in $\mathcal{C}$ is the product of f and g in $\mathcal{C}^{\rightarrow}$.*

Exercises 6.1.17 Sums. In a category $\mathcal{C}$, an operation $+$ is called a **sum** if, for every pair of objects A and B:

1. there are morphisms $i_{A,B} : A \rightarrow A + B$ and $j_{A,B} : B \rightarrow A + B$ called, respectively, the left and right **injections**
2. for every object C and any pair $f : A \rightarrow C$ and $g : B \rightarrow C$ of morphisms from A and B to C, there is a unique morphism $[f, g]_{A,B,C} : A + B \rightarrow C$ such that $f = [f, g]_{A,B,C} \circ i_{A,B}$ and $g = [f, g]_{A,B,C} \circ j_{A,B}$.

In other words, the following diagram commutes:

$$\begin{tikzcd}
A \arrow[r, "i"] \arrow[dr, "f"'] & A+B \arrow[d, dotted, "{[f,g]}"] & B \arrow[l, "j"'] \arrow[dl, "g"] \\
& C &
\end{tikzcd}$$

a) *The sum of two objects is unique up to isomorphism, when it exists.*

b) *The sum in $\mathcal{C}$ is the product in $\mathcal{C}^{op}$ (defined in 6.1.9).* This shows that sums and products enjoy dual properties.

Exponentials

We now interpret $\vdash_{\mathcal{N}}$ as the existence of morphisms, $\wedge$ by $\times$, and $\rightarrow$ by $\Rightarrow$

The $\rightarrow$-elimination rule shows that, whenever there are morphisms to both A and $A \Rightarrow B$, there must be a morphism to B too. The easiest way to ensure this is to take advantage of closure under composition and product, and simply require the existence of morphisms from $(A \Rightarrow B) \times A$ to B.

The $\rightarrow$-elimination rule shows that, whenever there is a morphism from $C \times A$ to B, there must be a morphism from C to $A \Rightarrow B$, too.

As already for 1 and $A \times B$, purely existential properties would not be enough to determine $A \Rightarrow B$ up to isomorphism, and thus we add a uniqueness condition.

Definition 6.1.18 *In a category with products*

$$\mathcal{C} = \langle Ob, Hom, \times \rangle,$$

an operation $\Rightarrow$ is called an **exponential** *if, for every triple of objects A, B and C:*

1. *there is a morphism $eval_{A,B} : (A \Rightarrow B) \times A \rightarrow B$ called* **evaluation**
2. *for every object C and any morphism $f : C \times A \rightarrow B$, there is a unique morphism $curry_{C,A,B}(f) : C \rightarrow (A \Rightarrow B)$ such that*

$$f = eval_{A,B} \circ (curry_{C,A,B}(f) \times id_A).$$

As usual, we can just say that the following diagram commutes:

$$\begin{tikzcd}
C \times A \arrow[r, "f"] \arrow[d, dotted, "curry(f) \times id"'] & B. \\
(A \Rightarrow B) \times A \arrow[ur, "eval"'] &
\end{tikzcd}$$

In terms of the deductive system introduced above, we have further axioms

$$\mathit{eval}_{A,B} : (A \Rightarrow B) \times A \to B,$$

and an additional rule

$$\frac{f : C \times A \to B}{\mathit{curry}\ (f) : C \to (A \Rightarrow B).}$$

While the latter literally corresponds to →-introduction, →-elimination can be stated as the derived rule

$$\frac{f : C \to (A \Rightarrow B) \qquad g : C \to A}{\mathit{eval}_{A,B} \circ (f \times g) : C \to B.}$$

Proposition 6.1.19 *The exponential of two objects is unique up to isomorphism, when it exists.*

Proof. Suppose $A \Rightarrow^* B$ also satisfies the properties, i.e.

1. there is a morphism $\mathit{eval}^* : (A \Rightarrow^* B) \times A \to B$
2. for every object C and any morphism $f : C \times A \to B$, there is a unique morphism $\mathit{curry}^*(f) : C \to (A \Rightarrow^* B)$ such that

$$f = \mathit{eval}^* \circ (\mathit{curry}^*(f) \times \mathit{id}_A).$$

Then, by letting $C = A \Rightarrow^* B$ and $f = \mathit{eval}^*$ in the definition of $A \Rightarrow B$, and $C = A \Rightarrow B$ and $f = \mathit{eval}$ in the definition of $A \Rightarrow^* B$, we have

$$\mathit{curry}\ (\mathit{eval}^*) \circ \mathit{curry}^*(\mathit{eval}\ ) = \mathit{id}_{A \Rightarrow B}$$

and

$$\mathit{curry}^*(\mathit{eval}\ ) \circ \mathit{curry}\ (\mathit{eval}^*) = \mathit{id}_{A \Rightarrow^* B}$$

by the uniqueness condition, and thus $A \Rightarrow B$ and $A \Rightarrow^* B$ are isomorphic. □

Exercise 6.1.20 The following hold in any category with products and exponentials.

a) $(\mathit{curry}\ f) \circ g = \mathit{curry}\ (f \circ (g \times \mathit{id}))$. (Hint: consider the following diagram.

$$\begin{tikzcd}
D \times A \arrow[r, "g \times id"] \arrow[rrd, "curry\ (f \circ (g \times id)) \times id"'] & C \times A \arrow[r, "f"] \arrow[rd, "(curry\ f) \times id"] & B \\
& & (A \Rightarrow B) \times A. \arrow[u, "eval"']
\end{tikzcd}$$

b) $\mathit{curry}\ (\mathit{eval}_{A,B}) = \mathit{id}_{A \Rightarrow B}$.

Exercises 6.1.21 The following hold in any category with terminal object, products and exponentials.

a) $A \Rightarrow (B \times C) \simeq (A \Rightarrow B) \times (A \Rightarrow C)$.

b) $A \Rightarrow (B \Rightarrow C) \simeq (A \times B) \Rightarrow C$.

c) $(1 \Rightarrow A) \simeq A$.

d) $(A \Rightarrow 1) \simeq 1$.

If we had written the exponential $A \Rightarrow B$ as the name implies, i.e. B^A, then the previous properties would become

$$(B \times C)^A \simeq B^A \times C^A \qquad (C^B)^A \simeq C^{A \times B} \qquad A^1 \simeq A \qquad \text{and} \qquad 1^A \simeq 1,$$

thus assuming a more familiar aspect.

Exercises 6.1.22 Given an abstract category with products $\mathcal{C}$ and two objects A and B, consider the structure $\mathcal{C}^{\Rightarrow}_{A,B}$ such that:

- the objects of $\mathcal{C}^{\Rightarrow}_{A,B}$ are the pairs (C, f), where C is an object of $\mathcal{C}$, and $f : C \times A \to B$ is a morphism of $\mathcal{C}$;
- the morphisms $h : (C_1, f_1) \to (C_2, f_2)$ of $\mathcal{C}^{\Rightarrow}_{A,B}$ are the morphisms $h : C_1 \to C_2$ of $\mathcal{C}$ such that $f_1 = f_2 \circ \langle h, id_A \rangle$:

$$\begin{tikzcd}
C_1 \times A \arrow[r, "f_1"] \arrow[d, dotted, "\langle h, id_A \rangle"'] & B. \\
C_2 \times A \arrow[ur, "f_2"'] &
\end{tikzcd}$$

a) *$\mathcal{C}^{\Rightarrow}_{A,B}$ is an abstract category.*

b) *The exponential $A \Rightarrow B$ in $\mathcal{C}$ is the terminal object in $\mathcal{C}^{\Rightarrow}_{A,B}$.* This gives an alternative proof of uniqueness of exponentials.

Cartesian closed categories

Having discovered the categorical properties forced on Hom, $\times$ and $\Rightarrow$ by the logical rules on $\vdash_{\mathcal{N}}$, $\wedge$ and $\rightarrow$, we now abstract them and introduce categories that, in the light of the previous discussion, turn out to be categorical models.

Definition 6.1.23 (Lawvere [1964]) *A* **cartesian closed category** *is a structure*

$$\mathcal{C} = \langle Ob, Hom, \times, \Rightarrow, 1 \rangle$$

such that:

1. $\langle Ob, Hom \rangle$ is a category

2. 1 *is the terminal object*

3. × *is the categorical product*

4. ⇒ *is the categorical exponential.*

The first three conditions define a category in which all finite families of objects have products (the terminal object is the product of the empty family).

Adjointness ⋆

The relationship between × and ⇒ expressed by the notion of exponential is reminiscent of the adjointness property in Heyting ⊓-algebras, and implies that the function

$$curry_{C,A,B} : Hom(C \times A, B) \to Hom(C, A \Rightarrow B)$$

is a set-theoretical isomorphism. Indeed:

- *one-onennes*

  If *curry* $(f) =$ *curry* (g), then

  $$f = eval \circ (curry(f) \times id_A) = eval \circ (curry(g) \times id_A) = g$$

  by definition of *curry*.

- *ontoness*

  Given any function $h : C \to (A \Rightarrow B)$, then

  $$h = curry(eval \circ (h \times id_A))$$

  by uniqueness of *curry*.

As in the case of Heyting ⊓-algebras, we could have defined the notion of exponential by first introducing a general notion of adjointness, and then requiring ⇒ to be the right adjoint of ×. The reason why we did not do so is that this road is more cumbersome and less intuitive, as the present subsection (which is not needed for the following) shows.

We introduce the notions of a 'function' on a category and an 'isomorphism' ≅ of *Hom* sets, so that we can generalize the condition

$$f(x) \sqsubseteq y \iff x \sqsubseteq g(y)$$

on partial orderings to

$$Hom(F(X), Y) \cong Hom(X, G(Y))$$

on categories.

The following is the appropriate notion of a 'function' on a category, as a map preserving all the categorical structure.

Definition 6.1.24 *A class function $F : \mathcal{C} \to \mathcal{C}$ is a* **functor** *on $\mathcal{C}$ if it associates an object $F(A)$ to each object A, and a morphism $F(f)$ to each morphism f, in such a way that:*

1. *if $f : A \to B$, then $F(f) : F(A) \to F(B)$*
2. *$F(id_A) = id_{F(A)}$*
3. *if $f : A \to B$ and $g : B \to C$, then $F(g \circ f) = F(g) \circ F(f)$.*

We can now introduce the notion of an 'identity' for functors on a category. Since objects in a category can only be identified up to isomorphism, two functors will be identical if their values are isomorphic (as sets). Moreover, we will ask that these values be uniformly isomorphic, in the sense of preserving the categorical structure.

Definition 6.1.25 *Two functors F and G on a category $\mathcal{C}$ are* **naturally isomorphic** *(written $\boldsymbol{F \cong G}$) if, for every object X of $\mathcal{C}$, there is a set-theoretical isomorphism (i.e. a one-one, onto function)*

$$\eta_X : F(X) \to G(X)$$

such that, for every morphism $f : X \to X'$,

$$G(f) \circ \eta_X = \eta_{X'} \circ F(f).$$

The definition expresses the fact that the following diagram commutes:

$$\begin{tikzcd}
F(X) \arrow[r, "\eta_X"] \arrow[d, "F(f)"] & G(X) \arrow[d, "G(f)"] \\
F(X') \arrow[r, "\eta_{X'}"] & G(X')
\end{tikzcd}$$

The notion of a natural isomorphism for functors is the inspiration for the following definition.[2]

Definition 6.1.26 (Kan [1958]) *A functor G on a category $\mathcal{C}$ is a* **right adjoint** *of a functor F on $\mathcal{C}$ (written $\mathbf{Hom}(\boldsymbol{F(X), Y}) \cong \mathbf{Hom}(\boldsymbol{X, G(Y)})$) if, for every pair of objects X and Y in $\mathcal{C}$, there is a set-theoretical isomorphism (i.e. a one-one, onto function)*

$$\eta_{X,Y} : Hom(F(X), Y) \to Hom(X, G(Y))$$

such that, for every triple $f : X' \to X$, $g : Y \to Y'$ and $h : F(X) \to Y$ of morphisms,

$$G(g) \circ \eta_{X,Y}(h) \circ f = \eta_{X',Y'}(g \circ h \circ F(f)).$$

[2] It is actually the same notion, applied to the functors from $\mathcal{C}^{op} \times \mathcal{C}$ to **Set** that map (X, Y) to $Hom(F(X), Y)$ and $Hom(X, G(Y))$, respectively. Here $\mathcal{C}^{op}$ is the opposite of $\mathcal{C}$ defined in 6.1.9, and **Set** is the category of all sets defined in 6.3.3.

The definition expresses the fact that the following diagram commutes:

$$\begin{tikzcd}
Hom(F(X), Y) \arrow[r, "\eta_{X,Y}"] \arrow[d] & Hom(X, G(Y)) \arrow[d] \\
Hom(F(X'), Y') \arrow[r, "\eta_{X',Y'}"] & Hom(X', G(Y')),
\end{tikzcd}$$

along the upper path

$$h \longmapsto \eta_{X,Y}(h) \longmapsto G(g) \circ \eta_{X,Y}(h) \circ f$$

and the lower path

$$h \longmapsto g \circ h \circ F(f) \longmapsto \eta_{X',Y}(g \circ h \circ F(f)).$$

Then, by letting $F_B(X) = X \times B$ and $G_B(Y) = B \Rightarrow Y$, G_B is a right adjoint of F_B, for every object B. We will abuse language, and say that $\Rightarrow$ *is a right adjoint of* $\times$.

The following fact extends 5.1.4.

Proposition 6.1.27 Uniqueness of Right Adjoints. *Given a category $\mathcal{C}$, a functor F on $\mathcal{C}$ has at most one right adjoint.*

Proof. We mimick the proof of 5.1.4, and discover during it the condition needed in the definition of functor and of natural isomorphism.

□

As for partial orderings, there is of course nothing special about *right* adjoints. In definition 6.1.26 we can say that F is a *left* adjoint of G, and a proof as above shows that left adjoints are unique, when they exist.

Exercise 6.1.28 *In any cartesian closed category, $Hom(X, Y) \cong Hom(1, X \Rightarrow Y)$.*

6.2 Soundness and Completeness Theorems

We can view cartesian closed categories as a further generalization of intuitionistic worlds. *Local truth* in a given cartesian closed category $\mathcal{C}$ under a given environment ρ now means that there is a morphism from 1 to $[\![\alpha]\!]_\rho$, and *global truth* means that there is such a morphism for every $\mathcal{C}$ and ρ. The following notion is then the analogue of intuitionistic and algebraic validity (2.2.4 and 5.2.1).

Definition 6.2.1 *A formula α is a* **categorical consequence** *of $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ (written $\boldsymbol{\Gamma \models_c \alpha}$) if, for every cartesian closed category $\mathcal{C}$ and every environment ρ on it,*

$$Hom([\![\gamma_1]\!]_\rho \times \cdots \times [\![\gamma_n]\!]_\rho, [\![\alpha]\!]_\rho) \neq \emptyset.$$

In the limit case of Γ empty, we get the notion of categorical validity: α is **categorically valid** (written $\models_c \alpha$) if $Hom(1, [\![\alpha]\!]_\rho) \neq \emptyset$ for every cartesian closed category, under every environment.

Theorem 6.2.2 Categorical Soundness. *For any* Γ *and* α*,*

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies \Gamma \models_c \alpha.$$

Proof. We prove that if $\mathcal{C}$ is a cartesian closed category, then it is a categorical model, i.e. if ρ is an environment on it and $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ and α are given, then

$$\Gamma \vdash_{\mathcal{N}} \alpha \implies Hom([\![\gamma_1]\!]_\rho \times \cdots \times [\![\gamma_n]\!]_\rho, [\![\alpha]\!]_\rho) \neq \emptyset.$$

The definition of cartesian closed category has been discovered precisely by looking at the properties that would make the present proof work, and thus we only have to repeat the work done above, proceeding by induction on $\vdash_{\mathcal{N}}$. For simplicity we write $[\![\Gamma]\!]_\rho$ for $[\![\gamma_1]\!]_\rho \times \cdots \times [\![\gamma_n]\!]_\rho$.

If $\Gamma, \beta \vdash_{\mathcal{N}} \beta$ is an assumption, then

$$Hom([\![\Gamma]\!]_\rho \times [\![\beta]\!]_\rho, [\![\beta]\!]_\rho) \neq \emptyset$$

follows from the existence of the right projection function $r_{[\![\Gamma]\!]_\rho, [\![\beta]\!]_\rho}$.

If $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ is obtained from $\Gamma, \alpha \vdash_{\mathcal{N}} \beta$ by →-introduction, then

$$Hom([\![\Gamma]\!]_\rho \times [\![\alpha]\!]_\rho, [\![\beta]\!]_\rho) \neq \emptyset$$

by the induction hypothesis, and so there is

$$f : [\![\Gamma]\!]_\rho \times [\![\alpha]\!]_\rho \to [\![\beta]\!]_\rho.$$

Then

$$curry\ (f) : [\![\Gamma]\!]_\rho \to ([\![\alpha]\!]_\rho \Rightarrow [\![\beta]\!]_\rho),$$

and so

$$Hom([\![\Gamma]\!]_\rho, [\![\alpha \to \beta]\!]_\rho) = Hom([\![\Gamma]\!]_\rho, [\![\alpha]\!]_\rho \Rightarrow [\![\beta]\!]_\rho) \neq \emptyset$$

by definition of $[\![\ \ ]\!]_\rho$.

If $\Gamma \vdash_{\mathcal{N}} \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \alpha \to \beta$ by →-elimination, then

$$Hom([\![\Gamma]\!]_\rho, [\![\alpha]\!]_\rho) \neq \emptyset$$

and

$$Hom([\![\Gamma]\!]_\rho, [\![\alpha]\!]_\rho \Rightarrow [\![\beta]\!]_\rho) = Hom([\![\Gamma]\!]_\rho, [\![\alpha \to \beta]\!]_\rho) \neq \emptyset$$

by definition of $[\![\ \ ]\!]_\rho$ and induction hypothesis, and so there are

$$f : [\![\Gamma]\!]_\rho \to ([\![\alpha]\!]_\rho \Rightarrow [\![\beta]\!]_\rho) \qquad \text{and} \qquad g : [\![\Gamma]\!]_\rho \to [\![\alpha]\!]_\rho.$$

Then

$$eval \circ \langle f, g\rangle : [\![\Gamma]\!]_\rho \to [\![\beta]\!]_\rho,$$

and so

$$Hom([\![\Gamma]\!]_\rho, [\![\beta]\!]_\rho) \neq \emptyset.$$

If $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ and $\Gamma \vdash_{\mathcal{N}} \beta$ by $\wedge$-introduction, then

$$Hom([\![\Gamma]\!]_\rho, [\![\alpha]\!]_\rho) \neq \emptyset \qquad \text{and} \qquad Hom([\![\Gamma]\!]_\rho, [\![\beta]\!]_\rho) \neq \emptyset$$

by the induction hypothesis, and so there are

$$f : [\![\Gamma]\!]_\rho \to [\![\alpha]\!]_\rho \qquad \text{and} \qquad g : [\![\Gamma]\!]_\rho \to [\![\beta]\!]_\rho.$$

Then

$$\langle f, g\rangle : [\![\Gamma]\!]_\rho \to [\![\alpha]\!]_\rho \times [\![\beta]\!]_\rho,$$

and so

$$Hom([\![\Gamma]\!]_\rho, [\![\alpha \wedge \beta]\!]_\rho) = Hom([\![\Gamma]\!]_\rho, [\![\alpha]\!]_\rho \times [\![\beta]\!]_\rho) \neq \emptyset$$

by definition of $[\![\ \ ]\!]_\rho$.

If $\Gamma \vdash_{\mathcal{N}} \alpha$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ by $\wedge$-elimination, then

$$Hom([\![\Gamma]\!]_\rho, [\![\alpha]\!]_\rho \times [\![\beta]\!]_\rho) = Hom([\![\Gamma]\!]_\rho, [\![\alpha \wedge \beta]\!]_\rho) \neq \emptyset$$

by definition of $[\![\ \ ]\!]_\rho$ and induction hypothesis, and so there is

$$f : [\![\Gamma]\!]_\rho \to [\![\alpha]\!]_\rho \times [\![\beta]\!]_\rho.$$

Then

$$l_{[\![\alpha]\!]_\rho, [\![\beta]\!]_\rho} \circ f : [\![\Gamma]\!]_\rho \to [\![\alpha]\!]_\rho,$$

and so

$$Hom([\![\Gamma]\!]_\rho, [\![\alpha]\!]_\rho) \neq \emptyset.$$

Similarly when $\Gamma \vdash_{\mathcal{N}} \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta$ by $\wedge$-elimination, this time using the right projection $r_{[\![\alpha]\!]_\rho, [\![\beta]\!]_\rho}$. □

We have noticed on p. 92 that the Algebraic Soundness Theorem is a strengthening of the Intuitionistic Soundness Theorem. We notice below (see 6.3.2) that, quite trivially, every Heyting ⊓-algebra is a cartesian closed category, from which it follows that *the Categorical Soundness Theorem is a strengthening of the Algebraic Soundness Theorem.* Thus cartesian closed categories provide the largest class of interpretations for Implicational Calculus with Conjunction introduced so far.

For the same reason, the following result is weaker than the Algebraic Completeness Theorem, which in turn was weaker than the Intuitionistic Completeness Theorem. We thus state it only for completeness, without proof.

Theorem 6.2.3 Categorical Completeness. *For any Γ and α,*

$$\Gamma \models_c \alpha \implies \Gamma \vdash_{\mathcal{N}} \alpha.$$

6.3 Examples $\star$

In this section we provide a series of examples of cartesian closed categories, selected with an eye to applications.

Heyting $\sqcap$-algebras

Our first example of a cartesian closed category connects the present notion with the one studied in the previous chapter.

Definition 6.3.1 *Given a Heyting $\sqcap$-algebra A, $\mathbf{A}$ is the categorical structure such that:*

- *$Ob_{\mathbf{A}}$ is the class of all singletons containing elements of A, i.e.*

$$Ob_{\mathbf{A}} = \{\{x\} : x \in A\};$$

- *for any pair of elements a and b of A,*

$$Hom_{\mathbf{A}}(\{a\}, \{b\}) \neq \emptyset \iff a \sqsubseteq b.$$

Notice that between $\{a\}$ and $\{b\}$ there is exactly one function, namely the one sending a to b: the previous definition says that such a function is a morphism of $\mathbf{A}$ exactly when $a \sqsubseteq b$.

Proposition 6.3.2 *For any Heyting $\sqcap$-algebra A, $\mathbf{A}$ is a cartesian closed category.*

Proof. $\mathbf{A}$ is obviously a category: the identities and composition exist because $\sqsubseteq$ is, respectively, reflexive and transitive. We thus concentrate on showing that $\mathbf{A}$ is cartesian closed.

- *terminal object*

  $\{ 1 \}$ is the terminal object. For any a, the existence of a (necessarily unique) morphism between $\{a\}$ and $\{1\}$ follows from the fact that 1 is the greatest element, i.e. that $a \sqsubseteq 1$.

- *products*

  Given two elements a and b, their categorical product is their g.l.b. $a \sqcap b$.

  The existence of the projections saying that $a \sqcap b \sqsubseteq a$ and $a \sqcap b \sqsubseteq b$ follows from the fact that $a \sqcap b$ is a *lower bound* of a and b.

  Moreover, given two morphisms f and g saying that $c \sqsubseteq a$ and $c \sqsubseteq b$, the existence of $\langle f, g \rangle$ saying that $c \sqsubseteq a \sqcap b$ follows from the fact that $a \sqcap b$ is the *greatest* lower bound of a and b.

- *exponentials*

  Given two elements a and b, their categorical exponential is $a \Rightarrow b$.

  The existence of *eval*, saying that $a \sqcap (a \Rightarrow b) \sqsubseteq b$, follows from 5.1.8.8 (which translates Modus Ponens).

  Moreover, given a morphism f saying that $c \sqcap a \sqsubseteq b$, the existence of *curry* (f) saying that $c \sqsubseteq (a \Rightarrow b)$ follows from the adjunction property of $\Rightarrow$ w.r.t. $\sqcap$. □

Sets

The previous examples of cartesian closed categories are quite trivial from the categorical point of view, since their *Hom* sets are severely restricted: they contain either *one* or *no* morphism. We now turn to the other end of the spectrum, and consider what is perhaps the most natural example of a cartesian closed category.

Definition 6.3.3 **Set** *is the categorical structure such that:*

- $Ob_{\mathbf{Set}}$ *is the class of all sets;*
- *for any pair of sets A and B, $Hom_{\mathbf{Set}}(A,B)$ is the set of all functions from A to B.*

Proposition 6.3.4 **Set** *is a cartesian closed category.*

Proof. Since **Set** is obviously a category, we concentrate on showing that it is cartesian closed. We define the appropriate objects and morphisms, and leave as an exercise the verification of the properties required by 6.1.23.

- *terminal object*

  We let 1 be any singleton set $\{x\}$ (up to isomorphism), and $t_A : A \to 1$ be the constant function with value x if A is non empty, and the undefined function if A is empty.

- *products*

  Given two sets A and B, we let

  $$A \times B = \{(x,y) : x \in A \ \wedge \ y \in B\}$$

  and

  $$l_{A,B}(x,y) = x \qquad \text{and} \qquad r_{A,B}(x,y) = y.$$

  Moreover, given two functions $f : C \to A$ and $g : C \to B$, we let

  $$\langle f, g\rangle(z) = (f(z), g(z)).$$

- *exponentials*

  Given two sets A and B, we let

  $$A \Rightarrow B = \{f : f \text{ is a function from } A \text{ to } B\},$$

  and

  $$eval_{A,B}(f,x) = f(x).$$

  Moreover, given a function $f : C \times A \to B$, we let

  $$(curry\ f)(z) = \text{the function } x \mapsto f(z,x). \quad \Box$$

Exercises 6.3.5 a) *Check that* 1*,* $A \times B$*, and* $A \Rightarrow B$ *satisfy the properties required by 6.1.23.*

b) *In* **Set***, a morphism is an isomorphism if and only if it is a one-one and onto function.* (Hint: it is enough to show that if $g \circ f = id$, then f is one-one and g is onto. For one-oneness, let $f(x) = f(y)$: then $g(f(x)) = g(f(y))$, and so $x = y$. For ontoness, given z, then $g(f(z)) = z$.)

c) *In* **Set***, two objects are isomorphic if and only if they are sets with the same cardinality.* (Hint: from part b).)

Chain complete partial orderings

The previous examples of cartesian closed categories are somewhat unsatisfactory, since they contain either too little or too much. We now want to exhibit a less trivial example, and experience with Heyting ⊓-algebras suggests us to look at the notion of topology.

The first thought would obviously be to look at the category **Top** having as objects the topological spaces, and as morphisms the continuous functions. However, this is not very satisfactory. First, the possibility of endowing any set with the *discrete topology*, in which every subset is open, shows that for many topological spaces continuity would not be a restriction. Second, the category **Top** is *not* cartesian closed anyway (see note 4 on p. 120).

We thus take an intermediate road and consider a notion of continuity that is, at the same time, sufficiently general but also sufficiently restricted.

Definition 6.3.6 (Markowsky [1976]) *A partially ordered set* $(D, \sqsubseteq_D)$ *is a* **chain complete partial ordering** *(***c.c.p.o.***) if every countable chain (i.e. every countable totally ordered set) of elements of* D *has a least upper bound (l.u.b.) in* D*.*

We will indicate by $\bigsqcup_{n\in\omega}^{D} x_n$ the l.u.b. of the chain $x_0 \sqsubseteq_D x_1 \sqsubseteq_D \cdots$. To increase readability we will omit reference to D and/or $\sqsubseteq_D$, when no confusion arises.

Notice that *every c.c.p.o. has a least element*, since we can consider in particular the empty chain: every element of D is (trivially) an upper bound to that chain, and the existence of the l.u.b. implies then the existence of the least element of D, which will be indicated by $\perp_D$.

A typical example of a c.c.p.o. is the set $\mathcal{P}(\omega)$ of all subsets of ω, ordered by inclusion, and with set-theoretical union as l.u.b.

We now introduce the appropriate notion of continuity.

Definition 6.3.7 *If $(D_1, \sqsubseteq_{D_1})$ and $(D_2, \sqsubseteq_{D_2})$ are two c.c.p.o.'s, then a function $f : D_1 \to D_2$ is:*

1. **monotone** *if it preserves the order, i.e.*

$$x \sqsubseteq_{D_1} y \;\Rightarrow\; f(x) \sqsubseteq_{D_2} f(y)$$

2. **chain continuous** *if it is monotone and it preserves l.u.b.'s of countable chains, i.e.*

$$f(\bigsqcup_{n\in\omega}^{D_1} x_n) = \bigsqcup_{n\in\omega}^{D_2} f(x_n).$$

Notice that the monotonicity condition in the definition of chain continuity is required to insure that the appropriate chains have l.u.b.'s,[3] as follows. If

$$x_0 \sqsubseteq_{D_1} x_1 \sqsubseteq_{D_1} \cdots$$

is a chain in D_1, then $\bigsqcup_{n\in\omega}^{D_1} x_n$ exists. If f is monotone,

$$f(x_0) \sqsubseteq_{D_2} f(x_1) \sqsubseteq_{D_2} \cdots$$

is still a chain in D_2, and $\bigsqcup_{n\in\omega}^{D_2} f(x_n)$ exists, and then it makes sense to require that

$$f(\bigsqcup_{n\in\omega}^{D_1} x_n) = \bigsqcup_{n\in\omega}^{D_2} f(x_n).$$

In general, *chain continuity is stronger than monotonicity*. For example, the function

$$f(A) = \begin{cases} \emptyset & \text{if } A \text{ is finite} \\ \omega & \text{otherwise} \end{cases}$$

is monotone on $\mathcal{P}(\omega)$. But it is not chain continuous because, if $A_n = \{0, 1, \ldots, n\}$, then $\bigcup_{n\in\omega} f(A_n)$ is $\emptyset$, while $f(\bigcup_{n\in\omega} A_n)$ is ω.

On the other hand, *on finite c.c.p.o.'s chain continuity coincides with monotonicity*: a monotone function always preserves l.u.b.'s of finite chains, and on a

[3] In other settings, continuity implies monotonicity: see e.g. 6.3.19.

finite c.c.p.o. every chain is finite. More generally, *the same holds for c.c.p.o.'s having no infinite ascending chain.*

The next exercises show that the word continuity has the usual topological meaning, in the sense that a function $f : A \to B$ on two topological spaces A and B is **continuous** if, for every open subset X of B, $f^{-1}(X) = \{x : f(x) \in X\}$ is an open subset of A.

Exercise 6.3.8 Order topology on a partial ordering. Recall from 5.3.8 that the **order topology** on a partial ordering is defined by taking the upward closed sets as the open sets.

The monotone functions on partial orderings are exactly the functions continuous w.r.t. their order topologies.

Exercises 6.3.9 Scott topology on a c.c.p.o. (Day and Kelly [1970], Scott [1972]) A subset U of a c.c.p.o. is **Scott open** if it is upward closed, and inaccessible by l.u.b.'s of chains. In other words, given any chain $x_0 \sqsubseteq x_1 \sqsubseteq \cdots$, if $(\bigsqcup_{n\in\omega} x_n) \in U$ then $x_n \in U$, for some n (and hence, by upward closure, for all sufficiently large n).

The topology defined by the Scott open sets is called the **Scott topology**, and the functions continuous w.r.t. it are called **Scott continuous**.

a) *The Scott open sets on a c.c.p.o. form a topology.* (Hint: unions and intersections of upward closed sets are still upward closed. If $(\bigsqcup_{n\in\omega} x_n) \in U_1 \cap U_2$, then x_n is in U_1 from a certain n_1 on, and in U_2 from a certain n_2 on. Then x_n is in $U_1 \cap U_2$ from $\max\{n_1, n_2\}$ on. If $\bigsqcup_{n\in\omega} x_n$ is in $\bigcup_{i\in I} U_i$, then it is in U_i for some i. Then x_n is in U_i, and hence in $\bigcup_{i\in I} U_i$, from a certain point on.)

b) *The chain continuous functions on a c.c.p.o. are exactly the Scott continuous functions.* (Hint: to show monotonicity of a Scott continuous function f, given x and y consider the Scott open set $X = \{z : z \not\sqsubseteq f(y)\}$, and use the fact that $f^{-1}(X)$ is Scott open, and hence upward closed, to show that if $f(x) \not\sqsubseteq f(y)$ then $x \not\sqsubseteq y$. To show preservation of the l.u.b. of a chain $\{x_n\}_{n\in\omega}$, consider the Scott open set $X = \{z : z \not\sqsubseteq \bigsqcup f(x_n)\}$, and use the fact that $f^{-1}(X)$ is Scott open, and hence inaccessible by l.u.b.'s of chains, to show that $f(\bigsqcup x_n) \sqsubseteq \bigsqcup f(x_n)$. The opposite direction follows from monotonicity.

To show Scott continuity of a chain continuous function f, consider a Scott open set X. Upward closure of $f^{-1}(X)$ follows from upward closure of X, using the fact that f is monotone. Inaccessibility by l.u.b.'s of chains of $f^{-1}(X)$ follows from inaccessibility of X, using the fact that f preserves l.u.b.'s of chains.)

c) *On $\mathcal{P}(\omega)$ the Scott open sets are exactly the unions of families $\{X : X \supseteq u\}$, where u is a finite set, and the Scott continuous functions are exactly the functions determined by their behavior on finite sets, in the sense that $f(X) = \bigcup_{u\subseteq X} f(u)$.* (Hint: if U is Scott open and $X \in U$, then $X = \bigcup_{n\in\omega} u_n$ for a chain $\{u_n\}_{n\in\omega}$ of finite sets. Then $u_n \in U$ for some n because U is inaccessible by l.u.b.'s of chains, and $\{X : X \supseteq u_n\} \subseteq U$ because U is upward closed. Conversely, any $U = \bigcup_{i\in I}\{X : X \supseteq u_i\}$ is obviously upward closed, and if $(\bigcup_{n\in\omega} X_n) \in U$ for some chain $\{X_n\}_{n\in\omega}$, then $(\bigcup_{n\in\omega} X_n) \supseteq u_i$ for some i, and $X_n \supseteq u_i$ for some n.)

Definition 6.3.10 Ccpo *is the categorical structure such that:*

- *$Ob_{\mathbf{Ccpo}}$ is the class of all c.c.p.o.'s;*
- *for any pair of c.c.p.o.'s D_1 and D_2, $Hom_{\mathbf{Ccpo}}(D_1, D_2)$ is the set of all chain continuous functions from D_1 to D_2.*

Before we turn to the proof that **Ccpo** is a cartesian closed category, we introduce the natural candidates for products and exponentials.

Definition 6.3.11 *If $(D_1, \sqsubseteq_{D_1})$ and $(D_2, \sqsubseteq_{D_2})$ are two c.c.p.o.'s,*

$$(\boldsymbol{D_1} \times \boldsymbol{D_2}, \sqsubseteq_{\boldsymbol{D_1} \times \boldsymbol{D_2}})$$

is defined as follows:

$$\begin{array}{rcl} \langle x, y\rangle \in D_1 \times D_2 & \Leftrightarrow & x \in D_1 \textit{ and } y \in D_2 \\ \langle x, y\rangle \sqsubseteq_{D_1 \times D_2} \langle x', y'\rangle & \Leftrightarrow & x \sqsubseteq_{D_1} x' \textit{ and } y \sqsubseteq_{D_2} y'. \end{array}$$

Proposition 6.3.12 *If D_1 and D_2 are two c.c.p.o.'s, then so is $D_1 \times D_2$.*

Proof. Let

$$\langle x_0, y_0\rangle \sqsubseteq_{D_1 \times D_2} \langle x_1, y_1\rangle \sqsubseteq_{D_1 \times D_2} \cdots$$

be a chain of elements of $D_1 \times D_2$. Then

$$\overset{D_1 \times D_2}{\bigsqcup_{n \in \omega}} \langle x_n, y_n\rangle = \langle \overset{D_1}{\bigsqcup_{n \in \omega}} x_n, \overset{D_2}{\bigsqcup_{n \in \omega}} y_n\rangle$$

because $\sqsubseteq_{D_1 \times D_2}$ is defined componentwise, using $\sqsubseteq_{D_1}$ and $\sqsubseteq_{D_2}$. □

To increase readability we will write $f(x, y)$ for $f(\langle x, y\rangle)$, when no confusion arises.

Given three c.c.p.o.'s D_i $(i = 1, 2, 3)$, a function

$$f : D_1 \times D_2 \to D_3,$$

is (by definition) *chain continuous on* $D_1 \times D_2$ if, whenever

$$\langle x_0, y_0\rangle \sqsubseteq_{D_1 \times D_2} \langle x_1, y_1\rangle \sqsubseteq_{D_1 \times D_2} \cdots,$$

then

$$f(\overset{D_1 \times D_2}{\bigsqcup_{n \in \omega}} \langle x_n, y_n\rangle) = \overset{D_3}{\bigsqcup_{n \in \omega}} f(x_n, y_n).$$

We say that f is *chain continuous on* D_1 alone if, whenever

$$x_0 \sqsubseteq_{D_1} x_1 \sqsubseteq_{D_1} \cdots,$$

then

$$f(\bigsqcup_{n\in\omega}^{D_1} x_n, y) = \bigsqcup_{n\in\omega}^{D_3} f(x_n, y).$$

Similarly, we say that f is *chain continuous on* D_2 alone if, whenever

$$y_0 \sqsubseteq_{D_2} y_1 \sqsubseteq_{D_2} \cdots,$$

then

$$f(x, \bigsqcup_{n\in\omega}^{D_2} y_n) = \bigsqcup_{n\in\omega}^{D_3} f(x, y_n).$$

We now prove that these notions of chain continuity are nicely related, in the following sense.

Proposition 6.3.13 *A function f is chain continuous on $D_1 \times D_2$ if and only if it is chain continuous on D_1 and D_2 separately.*[4]

Proof. If f is chain continuous on $D_1 \times D_2$ and

$$x_0 \sqsubseteq_{D_1} x_1 \sqsubseteq_{D_1} \cdots,$$

then

$$\langle x_0, y\rangle \sqsubseteq_{D_1\times D_2} \langle x_1, y\rangle \sqsubseteq_{D_1\times D_2} \cdots$$

by definition of $\sqsubseteq_{D_1\times D_2}$, and hence

$$f(\bigsqcup_{n\in\omega}^{D_1} x_n, y) = f(\bigsqcup_{n\in\omega}^{D_1\times D_2} \langle x_n, y\rangle) = \bigsqcup_{n\in\omega}^{D_3} f(x_n, y)$$

by definition of $\bigsqcup_{n\in\omega}^{D_1\times D_2}$, and chain continuity of f on $D_1 \times D_2$. Thus f is chain continuous on D_1. Similarly, f is chain continuous on D_2.

Conversely, suppose f is chain continuous on D_1 and D_2 separately, and

$$\langle x_0, y_0\rangle \sqsubseteq_{D_1\times D_2} \langle x_1, y_1\rangle \sqsubseteq_{D_1\times D_2} \cdots.$$

[4] Notice that this property fails for the usual notion of continuity on the real numbers. For example, the function

$$f(x,y) = \begin{cases} \frac{xy}{x^2+y^2} & \text{if } x \neq 0 \text{ or } y \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

is continuous at $(0,0)$ in both x and y separately, but not in (x,y) together (Peano [1884]). Indeed, $f(x,y) = 0$ if exactly one of x and y is equal to 0, but $f(x,y) = \frac{1}{2}$ for $x = y$ and $x, y \neq 0$.

Then

$$\begin{array}{rcl} f(\bigsqcup_{n\in\omega}^{D_1\times D_2}\langle x_n, y_n\rangle) & = & f(\bigsqcup_{n\in\omega}^{D_1} x_n, \bigsqcup_{n\in\omega}^{D_2} y_n) \\ & = & \bigsqcup_{p\in\omega}^{D_3} f(x_p, \bigsqcup_{n\in\omega}^{D_2} y_n) \\ & = & \bigsqcup_{p\in\omega}^{D_3}\bigsqcup_{q\in\omega}^{D_3} f(x_p, y_q) \\ & = & \bigsqcup_{n\in\omega}^{D_3} f(x_n, y_n) \end{array}$$

by definition of $\bigsqcup_{n\in\omega}^{D_1\times D_2}$, chain continuity of f on D_1, chain continuity of f on D_2, and monotonicity of f, which implies

$$f(x_p, y_q) \sqsubseteq_{D_3} f(x_n, y_n)$$

for any $n \geq p, q$. $\square$

We can now turn to the other fundamental notion of cartesian closed categories.

Definition 6.3.14 *If $(D_1, \sqsubseteq_{D_1})$ and $(D_2, \sqsubseteq_{D_2})$ are two c.c.p.o.'s,*

$$(\boldsymbol{[D_1 \to D_2]}, \sqsubseteq_{\boldsymbol{[D_1\to D_2]}})$$

is defined as follows:

$$\begin{array}{rcl} f \in [D_1 \to D_2] & \Leftrightarrow & f \textit{ is a chain continuous function from } D_1 \textit{ to } D_2 \\ f \sqsubseteq_{[D_1\to D_2]} g & \Leftrightarrow & (\forall x \in D_1)(f(x) \sqsubseteq_{D_2} g(x)). \end{array}$$

In other words, chain continuous functions from D_1 to D_2 are naturally ordered by their graphs, pointwise.

Proposition 6.3.15 *If D_1 and D_2 are two c.c.p.o.'s, then so is $[D_1 \to D_2]$.*

Proof. Let

$$f_0 \sqsubseteq_{[D_1\to D_2]} f_1 \sqsubseteq_{[D_1\to D_2]} \cdots$$

be a chain of functions in $[D_1 \to D_2]$. By definition, for every $x \in D_1$,

$$f_0(x) \sqsubseteq_{D_2} f_1(x) \sqsubseteq_{D_2} \cdots.$$

We can thus define the l.u.b. of $\{f_n\}_{n\in\omega}$ as the function $\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n$ that behaves as follows, for any $x \in D_1$:

$$(\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n)(x) = \bigsqcup_{n\in\omega}^{D_2} f_n(x).$$

Then:

1. $\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n$ *is a chain continuous function from* D_1 *to* D_2

   If $x_0 \sqsubseteq_{D_1} x_1 \sqsubseteq_{D_1} \cdots$, then

$$\begin{aligned}(\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n)(\bigsqcup_{m\in\omega}^{D_1} x_m) &= \bigsqcup_{n\in\omega}^{D_2}(f_n(\bigsqcup_{m\in\omega}^{D_1} x_m)) \\ &= \bigsqcup_{n\in\omega}^{D_2}(\bigsqcup_{m\in\omega}^{D_2} f_n(x_m)) \\ &= \bigsqcup_{m\in\omega}^{D_2}(\bigsqcup_{n\in\omega}^{D_2} f_n(x_m)) \\ &= \bigsqcup_{m\in\omega}^{D_2}(\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n)(x_m)\end{aligned}$$

   by definition of l.u.b. in $[D_1 \to D_2]$, chain continuity of f_n, commutativity of l.u.b.'s (which can be proved as an exercise), and definition of l.u.b. in $[D_1 \to D_2]$ again.

2. $\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n$ *is the l.u.b. of* $\{f_n\}_{n\in\omega}$ *in* $[D_1 \to D_2]$

   $\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n$ is an upper bound by definition of $\sqsubseteq_{[D_1\to D_2]}$ because, for every $x \in D_1$,

$$f_m(x) \sqsubseteq_{D_2} \bigsqcup_{n\in\omega}^{D_2} f_n(x) = (\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n)(x),$$

   and so $f_m \sqsubseteq_{[D_1\to D_2]} (\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n)$, for every m.

   Suppose g is any other upper bound, i.e.

$$f_n \sqsubseteq_{[D_1\to D_2]} g,$$

   for every n. Then

$$f_n(x) \sqsubseteq_{D_2} g(x)$$

for every $x \in D_1$, and

$$(\bigsqcup_{n\in\omega}^{[D_1\to D_2]} f_n)(x) = \bigsqcup_{n\in\omega}^{D_2} f_n(x) \sqsubseteq_{D_2} g(x)$$

by definition of $\sqsubseteq_{[D_1\to D_2]}$, and properties of l.u.b.'s. □

Notice that $\bot_{[D_1\to D_2]}$, which must exist because $[D_1 \to D_2]$ is a c.c.p.o., is the constant function on D_1 with value $\bot_{D_2}$.

We now have all the ingredients needed to prove the main result.

Proposition 6.3.16 **Ccpo** *is a cartesian closed category.*

Proof. **Ccpo** is a category: the identities are obviously chain continuous functions, and composition is naturally defined as usual. We thus concentrate on showing that **Ccpo** is cartesian closed.

- *terminal object*

  We show that the c.c.p.o. $\{\bot\}$ is terminal. If D is any c.c.p.o., then there is exactly one function $f : D \to \{\bot\}$, namely the constant function with value $\bot$ if D is nonempty, and the undefined function if D is empty. Such a function is obviously chain continuous.

- *products*

  Given two c.c.p.o.'s $(D_1, \sqsubseteq_1)$ and $(D_2, \sqsubseteq_2)$, we show that their categorical product is $(D_1 \times D_2, \sqsubseteq_{D_1\times D_2})$.

  We obviously define

  $$l_{D_1,D_2}(x, y) = x \qquad \text{and} \qquad r_{D_1,D_2}(x, y) = y$$

  and, given two functions $f : D \to D_1$ and $g : D \to D_2$,

  $$\langle f, g\rangle(z) = \langle f(z), g(z)\rangle.$$

  These functions satisfy the properties required by the definition of product, since they already do so in **Set**. It only remains to show that they are morphisms of **Ccpo**, i.e. that they are chain continuous.

  Chain continuity of the projections follows immediately from 6.3.13 and the fact that both the identities and the constant functions are chain continuous.

  For the chain continuity of $\langle f, g\rangle$, let f and g be chain continuous and

  $$x_0 \sqsubseteq_D x_1 \sqsubseteq_D \cdots.$$

Then:

$$\begin{array}{rcl}\langle f,g\rangle(\bigsqcup\limits_{n\in\omega}^{D} x_n) & = & \langle f(\bigsqcup\limits_{n\in\omega}^{D} x_n), g(\bigsqcup\limits_{n\in\omega}^{D} x_n)\rangle \\ & = & \langle \bigsqcup\limits_{n\in\omega}^{D_1} f(x_n), \bigsqcup\limits_{n\in\omega}^{D_2} g(x_n)\rangle \\ & = & \bigsqcup\limits_{n\in\omega}^{D_1\times D_2} \langle f(x_n), g(x_n)\rangle \\ & = & \bigsqcup\limits_{n\in\omega}^{D_1\times D_2} \langle f,g\rangle(x_n)\end{array}$$

by definition of $\langle f,g\rangle$, chain continuity of f and g, definition of l.u.b. in $D_1 \times D_2$, and definition of $\langle f,g\rangle$ again.

- *exponentials*

  Given two c.c.p.o.'s $(D_1, \sqsubseteq_1)$ and $(D_2, \sqsubseteq_2)$, we show that their categorical exponential is $([D_1 \to D_2], \sqsubseteq_{[D_1\to D_2]})$.

  We obviously define

  $$eval_{D_1,D_2}(f,x) = f(x)$$

  and, given a function $f : D \times D_1 \to D_2$,

  $$(curry\ f)(z) = \text{the function } x \mapsto f(z,x).$$

  These functions satisfy the properties required by the definition of exponential, since they already do so in **Set**. It only remains to show that they are morphisms of **Ccpo**, i.e. that they are chain continuous.

  For *eval*, by 6.3.13 it is enough to show chain continuity in each variable separately.

  1. *first variable*

     If

     $$f_0 \sqsubseteq_{[D_1\to D_2]} f_1 \sqsubseteq_{[D_1\to D_2]} \cdots,$$

     then, for any $x \in D_1$,

     $$\begin{array}{rcl}eval_{D_1,D_2}(\bigsqcup\limits_{n\in\omega}^{[D_1\to D_2]} f_n, x) & = & (\bigsqcup\limits_{n\in\omega}^{[D_1\to D_2]} f_n)(x) \\ & = & \bigsqcup\limits_{n\in\omega}^{D_2} f_n(x)\end{array}$$

$$= \bigsqcup_{n\in\omega}^{D_2} eval_{D_1,D_2}(f_n, x)$$

by definition of $eval_{D_1,D_2}$, of $\bigsqcup_{n\in\omega}^{[D_1\to D_2]}$, and of $eval_{D_1,D_2}$ again.

2. *second variable*
   If
   $$x_0 \sqsubseteq_{D_1} x_1 \sqsubseteq_{D_1} \cdots,$$
   then, for any $f \in [D_1 \to D_2]$,
   $$\begin{array}{rcl} eval_{D_1,D_2}(f, \bigsqcup_{n\in\omega}^{D_1} x_n) & = & f(\bigsqcup_{n\in\omega}^{D_1} x_n) \\ & = & \bigsqcup_{n\in\omega}^{D_2} f(x_n) \\ & = & \bigsqcup_{n\in\omega}^{D_2} eval_{D_1,D_2}(f, x_n) \end{array}$$
   by definition of $eval_{D_1,D_2}$, chain continuity of f, and definition of $eval_{D_1,D_2}$ again.

For *curry*, we need to show two things:

1. *curry f is well-defined*
   This amounts to show that, for every $z \in D$, $(curry\ f)(z) \in [D_1 \to D_2]$, i.e. that $(curry\ f)(z)$ is a chain continuous function from D_1 to D_2, and it follows from 6.3.13: since f is chain continuous as a function of two variables, it is also chain continuous as a function of the second variable separately (with z fixed).
2. *curry f is chain continuous*
   If
   $$z_0 \sqsubseteq_D z_1 \sqsubseteq_D \cdots,$$
   then
   $$\begin{array}{rcl} (curry\ f)(\bigsqcup_{n\in\omega}^{D} z_n) & = & \text{the function } x \mapsto f(\bigsqcup_{n\in\omega}^{D} z_n, x) \\ & = & \text{the function } x \mapsto \bigsqcup_{n\in\omega}^{D} f(z_n, x) \\ & = & \bigsqcup_{n\in\omega}^{[D_1\to D_2]} (\text{functions } x \mapsto f(z_n, x)) \end{array}$$

$$= \bigsqcup_{n\in\omega}^{[D_1\to D_2]} (\mathit{curry}\ f)(z_n)$$

by definition of *curry*, chain continuity of f, definition of $\bigsqcup_{n\in\omega}^{[D_1\to D_2]}$, and definition of *curry* again. □

Complete partial orderings

While **Ccpo** is a nice restriction of **Set**, we can do better by imposing further closure conditions.

Definition 6.3.17 *A partially ordered set* $(D, \sqsubseteq_D)$ *is a* **complete partial ordering (c.p.o.)** *if every set of elements of* D *has a least upper bound in* D.

Notice that *a complete partial ordering is a complete lattice*, since the l.u.b. is automatically defined, and the g.l.b. is the l.u.b. of the lower bounds.

Obviously, *every c.p.o. is a c.c.p.o.* (in particular, it has a least element $\perp$), *but not conversely*. A counterexample is provided by the c.c.p.o. used in 6.3.26, which is not a c.p.o. because the set $\{y, z\}$ has no l.u.b. A typical example of a c.p.o. is the set $\mathcal{P}(\omega)$ already considered as an example of a c.c.p.o.

The notion of monotonicity obviously applies to c.p.o.'s. We now extend the notion of chain continuity.

Definition 6.3.18 *If* D_1 *and* D_2 *are two c.p.o.'s, a function* $f : D_1 \to D_2$ *is* **continuous** *if it preserves arbitrary l.u.b.'s, i.e. for any subset* X *of* D_1

$$f(\bigsqcup^{D_1} X) = \bigsqcup^{D_2} f(X),$$

where $f(X) = \{f(x) : x \in X\}$.

The reason we did not require monotonicity in the definition of continuity, as for chain continuity in 6.3.7, is the following.

Proposition 6.3.19 *A continuous function on a c.p.o. is monotone.*

Proof. If f is continuous, then

$$\begin{aligned} x \sqsubseteq y &\implies y = \bigsqcup\{x, y\} \\ &\implies f(y) = \bigsqcup\{f(x), f(y)\} \\ &\implies f(x) \sqsubseteq f(y). \quad \square \end{aligned}$$

Obviously, the crucial fact in the previous proof is that l.u.b.'s always exist in c.p.o.'s, and in particular they exist for the pair of elements $f(x)$ and $f(y)$. On c.c.p.o.'s they would instead exist, in general, only if $f(x)$ and $f(y)$ form a chain, which is the case if f is monotone. The previous argument thus fails for arbitrary c.c.p.o.'s.

We can now introduce the announced restriction of **Ccpo**.

Definition 6.3.20 Cpo *is the categorical structure such that:*

- *$Ob_{\mathbf{Cpo}}$ is the class of all c.p.o.'s;*
- *for any pair of c.p.o.'s D_1 and D_2, $Hom_{\mathbf{Cpo}}(D_1, D_2)$ is the set of all continuous functions from D_1 to D_2.*

We leave to the reader the verification that all properties proved in the present subsection for c.c.p.o.'s and chain continuous functions continue to hold for c.p.o.'s and continuous functions. In particular:

Proposition 6.3.21 Cpo *is a cartesian closed category.*

Exercises 6.3.22 Directed or bounded sets. The notions of c.c.p.o. and c.p.o. are obtained by requiring the existence of l.u.b.'s for countable chains and arbitrary subsets, respectively. Intermediate between these two extremes are other possibilies, in particular requiring the existence of l.u.b.'s for:

- **directed subsets** X: for any pair of elements $x_1, x_2 \in X$ there exists an element $y \in X$ such that $x_1, x_2 \sqsubseteq y$;
- **bounded subsets** X: there exists an element y such that, for every $x \in X$, $x \sqsubseteq y$ (i.e. X has an upper bound).

The notions of **directed complete partial ordering** (d.c.p.o.) and **bounded complete partial ordering** (b.c.p.o.) are obtained by requiring the existence of l.u.b.'s for all directed and all bounded subsets, respectively. The notions of **directed continuous function** and **bounded continuous function** are defined as is 6.3.18, by requiring monotonicity and preservation of l.u.b.'s of directed and bounded subsets, respectively. The categorical structures **Dcpo** and **Bcpo** are obtained in the obvious way.

a) *The following are all the implications among the notions of completeness introduced so far:*

$$\begin{tikzcd}
\textit{c.p.o.} \arrow[r] \arrow[dr] & \textit{d.c.p.o.} \arrow[r] & \textit{c.c.p.o.} \\
 & \textit{b.c.p.o.} &
\end{tikzcd}$$

(Hint: for the positive part, it is enough to note that a chain is a directed set. For the negative part: the set $\mathcal{P}$ of partial functions on ω ordered by their graphs is an example of a d.c.p.o. that is not complete; the set ω_1 of countable ordinals ordered by magnitude is an example of a c.c.p.o. that is not directed complete; the set ω of finite ordinals ordered by magnitude is an example of a b.c.p.o. that is not chain complete; and the following

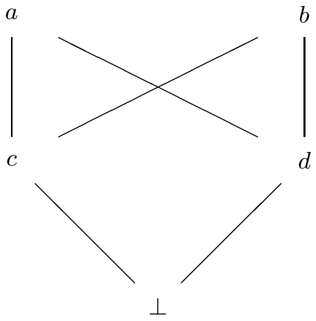

is an example of a d.c.p.o. that is not bounded complete.)

b) **Dcpo** *and* **Bcpo** *are cartesian closed categories.*

Algebraic partial orderings

Having defined an abstract notion of continuity for (chain) complete partial orderings, we now attempt a similar abstraction for the notion of finiteness. The starting point is given by the following trivial but crucial property.

Proposition 6.3.23 *On the c.c.p.o. $\mathcal{P}(\omega)$ of the sets of natural numbers ordered by inclusion, a set A is finite if and only if, whenever $A \subseteq \bigcup_{n\in\omega} X_n$ for a chain*

$$X_0 \subseteq X_1 \subseteq \cdots,$$

then $A \subseteq X_n$ for some n.

Proof. If A is finite, let $\{a_0, \ldots, a_m\}$ be an enumeration of its elements. If

$$X_0 \subseteq X_1 \subseteq \cdots$$

and $A \subseteq \bigcup_{n\in\omega} X_n$, for every $i \leq m$ there is n_i such that $a_i \in X_{n_i}$. Since the chain is increasing, $A \subseteq X_{\max\{n_0,\ldots,n_m\}}$.

If A is infinite, let $\{a_0, a_1, \ldots\}$ be an enumeration of its elements. Consider the increasing chain of sets defined by

$$X_n = \{a_0, \ldots, a_n\}.$$

Then $A = \bigcup_{n\in\omega} X_n$, but $A \not\subseteq X_n$ for any n, since X_n is finite and A is not. □

The characterization of finiteness just given is purely order theoretic, and it can thus be extended to arbitrary c.c.p.o.'s.

Definition 6.3.24 *In a c.c.p.o. an element a is called* **finite** *if, whenever $a \sqsubseteq \bigsqcup x_n$ for a chain $x_0 \sqsubseteq_D x_1 \sqsubseteq_D \cdots$, then $a \sqsubseteq x_n$ for some n.*

A c.c.p.o. is called **algebraic** *if, for every element x, there is a chain of finite elements $x_0 \sqsubseteq_D x_1 \sqsubseteq_D \cdots$ such that $x = \bigsqcup x_n$. In other words, the finite elements generate the c.c.p.o.*

Exercises 6.3.25 a) *The algebraic c.c.p.o.'s, with the chain continuous functions as morphisms, form a category.*

b) *Every finite c.c.p.o. is algebraic.*

c) *If D_1 and D_2 are algebraic c.c.p.o.'s, then so is $D_1 \times D_2$.* (Hint: the finite elements of $D_1 \times D_2$ are exactly the pairs of finite elements of D_1 and D_2.)

Despite the good news of the previous exercises, the next result shows that the notion of algebraic c.c.p.o. is not adequate for our purposes.

Proposition 6.3.26 *There is an algebraic c.c.p.o. D such that $[D \to D]$ is not algebraic.*

Proof. Let

$$D = \{\bot, y, z, x_0, \dots, x_n, \dots\}$$

with the ordering described by the next diagram:

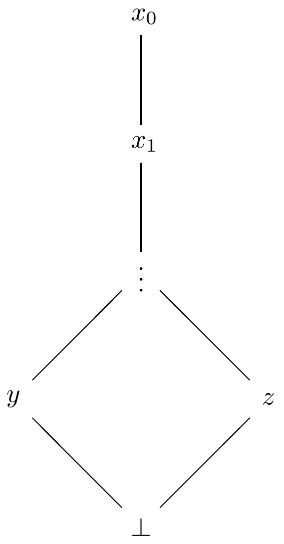

D is obviously an algebraic c.c.p.o., because there is no infinite ascending chain. Thus every element is finite.

To prove that $[D \to D]$ is not algebraic, we show that the identity function id_D is not the l.u.b. of the finite elements below it. Since, obviously, id_D cannot be the l.u.b. of functions with range contained in $\{\bot, y, z\}$, it will be enough to show that no monotone function f such that $f \sqsubseteq_{[D \to D]} id_D$ and with range not contained in $\{\bot, y, z\}$ is finite.

Let thus $f \sqsubseteq_{[D \to D]} id_D$, i.e. $f(x) \sqsubseteq_D x$ for every x. In particular, $f(y) \sqsubseteq_D y$ and $f(z) \sqsubseteq z$. Since the range of f is not contained in $\{\bot, y, z\}$, there is some n_0 such that $f(x_{n_0}) \sqsupseteq_D y, z$. Then, by monotonicity, $f(x_n) \sqsupseteq_D y, z$ for all $n \geq n_0$. We can thus define, for each $n \geq n_0$, functions f_n as follows:

$$f_n(x) = \begin{cases} f(x) & \text{if } x = \bot, y, z, x_0, \dots, x_n \\ \text{the predecessor of } f(x) \text{ w.r.t. } \sqsubseteq_D & \text{otherwise.} \end{cases}$$

Then:

- *for any $n \geq n_0$, f_n is chain continuous*

  It is obviously monotone, and since there is no infinite ascending chain in D, this is equivalent to being chain continuous.

- *for any $n \geq n_0$, $f_n \sqsubseteq f_{n+1}$*

  For any $x \neq x_{n+1}$, $f_n(x) = f_{n+1}(x)$. And for $x = x_{n+1}$,

  $$f_n(x_{n+1}) = \text{the predecessor of } f(x_{n+1}) \sqsubset f(x_{n+1}) = f_{n+1}(x_{n+1}).$$

- $f = \bigsqcup_{n \geq n_0} f_n$

  For $x = \bot, y, z$, $f_n(x) = f(x)$ for any n. For $x = x_m$, $f_n(x) = f(x)$ for any $n \geq m$.

- *for any $n \geq n_0$, $f \not\sqsubseteq f_n$*

  It is enough to show that, for some x, $f(x) \not\sqsubseteq f_n(x)$. For $x = x_{n+1}$,

  $$f(x_{n+1}) \not\sqsubseteq f_n(x_{n+1}) = \text{the predecessor of } f(x_{n+1}).$$

Thus f is not finite. □

The previous example used in a crucial way the existence of sets (actually, pairs) of elements without l.u.b. Since the set D is not only a c.c.p.o. but also a d.c.p.o. (because any directed set of elements contains some x_n, and hence it has a l.u.b.), we turn to the consideration of c.p.o.'s.[5]

The starting point is again given by a simple fact on sets of natural numbers.

[5] Although the notions of directed *or* bounded completeness are not enough individually, their combination (directed *and* bounded completeness) would be.

Proposition 6.3.27 *On the c.p.o. $\mathcal{P}(\omega)$ a nonempty set A is a singleton if and only if, whenever $A \subseteq \bigcup \mathcal{X}$ for a set $\mathcal{X} \subseteq \mathcal{P}(\omega)$, then $A \subseteq X$ for some $X \in \mathcal{X}$.*

Proof. If A is a singleton $\{a\}$ and $A \subseteq \bigcup \mathcal{X}$, then $a \in \bigcup \mathcal{X}$. By definition of union, $a \in X$ for some $X \in \mathcal{X}$, i.e. $A \subseteq X$.

If A is not a singleton, consider the family

$$\mathcal{X} = \{\{a\} : a \in A\}.$$

Then $A \subseteq \bigcup \mathcal{X}$. But $A \not\subseteq \{a\}$ for any $a \in A$, since A is not a singleton. □

The characterization just given of being a singleton, which is a strong version of finiteness, is again purely order theoretic, and can thus be extended to arbitrary c.p.o.'s.

Definition 6.3.28 *In a c.p.o. an element a is called* **strongly finite** *if, whenever $a \sqsubseteq \bigsqcup X$ for a subset X, then $a \sqsubseteq x$ for some $x \in X$.*

A c.p.o. is called **algebraic** *if every element x is the l.u.b. of the strongly finite elements below it, i.e.*

$$x = \bigsqcup \{a : a \sqsubseteq x \text{ and } a \text{ strongly finite}\}.$$

Definition 6.3.29 **Alg** *is the categorical structure such that:*

- *$Ob_{\mathbf{Alg}}$ is the class of all algebraic c.p.o.'s;*
- *for any pair of algebraic c.c.p.o.'s D_1 and D_2, $Hom_{\mathbf{Alg}}(D_1, D_2)$ is the set of all continuous functions from D_1 to D_2.*

Proposition 6.3.30 **Alg** *is a cartesian closed category.*

Proof. Since **Cpo** is a cartesian closed category, it is enough to show that **Alg** is closed with respect to the appropriate operations.

- *terminal object*

  The c.p.o. $\{\bot\}$, which is the terminal object in **Cpo**, is algebraic because it is finite.

- *products*

  That $D_1 \times D_2$ is an algebraic c.p.o. if D_1 and D_2 are, follows from the fact that the strongly finite elements of $D_1 \times D_2$ are exactly the pairs of storngly finite elements of D_1 and D_2.

  Let a and b be strongly finite in D_1 and D_2, and $\langle a, b\rangle \sqsubseteq_{D_1 \times D_2} \bigsqcup X$. If

  $$X_1 = \{x_1 : \langle x_1, x_2\rangle \in X\} \qquad \text{and} \qquad X_2 = \{x_2 : \langle x_1, x_2\rangle \in X\},$$

then $a \sqsubseteq_{D_1} \bigsqcup X_1$ and $b \sqsubseteq_{D_2} \bigsqcup X_2$. By strong finiteness, $a \sqsubseteq x_1$ and $b \sqsubseteq x_2$ for some $x_1 \in X_1$ and $x_2 \in X_2$, i.e. $\langle a, b\rangle \sqsubseteq \langle x_1, x_2\rangle \in X$.

Let now $\langle a, b\rangle$ be strongly finite in $D_1 \times D_2$. To show that a is strongly finite in D_1, let $a \sqsubseteq \bigsqcup X_1$. Then $\langle a, b\rangle \sqsubseteq \bigsqcup(X_1 \times \{b\})$, and by strong finiteness $\langle a, b\rangle \sqsubseteq \langle x_1, b\rangle$ for some $x_1 \in X_1$, i.e. $a \sqsubseteq x_1$. Similarly, b is strongly finite in D_2.

- *exponentials*

  We prove that $[D_1 \to D_2]$ is an algebraic c.p.o., if D_1 and D_2 are. For any pair of elements a of D_1 and b of D_2, consider the following step function:

  $$f_{ab}(x) = \begin{cases} b & \text{if } a \sqsubseteq x \\ \bot_{D_2} & \text{otherwise.} \end{cases}$$

  We first note the following:

  - f_{ab} *is continuous, i.e. it belongs to* $[D_1 \to D_2]$

    Given any set X, we consider two exhaustive cases.

    If $a \sqsubseteq \bigsqcup X$, then $f(\bigsqcup X) = b$ by definition. Moreover, by strong finiteness, $a \sqsubseteq x$ for some $x \in X$, and hence also $\bigsqcup f(X) = b$.

    If $a \not\sqsubseteq \bigsqcup X$, then $f(\bigsqcup X) = \bot$ by definition. Moreover, for every $x \in X$ it must also be $a \not\sqsubseteq x$, and thus $\bigsqcup f(X) = \bot$.

    In both cases, then, $f(\bigsqcup X) = \bigsqcup f(X)$.

  - *if a is strongly finite in D_1 and b is strongly finite in D_2, then f_{ab} is strongly finite in* $[D_1 \to D_2]$

    Given any set F of continuous functions, suppose $f_{ab} \sqsubseteq_{[D_1 \to D_2]} \bigsqcup F$. We want to find $f \in F$ such that $f_{ab} \sqsubseteq_{[D_1 \to D_2]} f$, i.e. $f_{ab}(x) \sqsubseteq_{D_2} f(x)$ for every $x \in D_1$. By definition of f_{ab}, it is enough to find $f \in F$ such that $b \sqsubseteq_{D_2} f(x)$, for every $x \in D_1$ such that $a \sqsubseteq x$. But

    $$b = f_{ab}(a) \sqsubseteq_{D_2} (\overset{[D_1 \to D_2]}{\bigsqcup} F)(a) = \overset{D_2}{\bigsqcup}\{f(a) : f \in F\}.$$

    Since b is strongly finite, there is $f \in F$ such that $b \sqsubseteq_{D_2} f(a)$. Moreover, if $a \sqsubseteq x$, then $f(a) \sqsubseteq f(x)$ because f is monotone, and hence $b \sqsubseteq f(x)$ for all such x.

  To prove that $[D_1 \to D_2]$ is algebraic it is now enough to show that, for every continuous $f : D_1 \to D_2$,

  $$f = \overset{[D_1 \to D_2]}{\bigsqcup} \{f_{ab} : f_{ab} \sqsubseteq_{[D_1 \to D_2]} f \ \wedge \ f_{ab} \text{ is strongly finite}\}.$$

Actually, we prove that

$$f = \bigsqcup^{[D_1 \to D_2]} \{f_{ab} : b \sqsubseteq f(a) \ \wedge \ f_{ab} \text{ is strongly finite}\}.$$

- $(\bigsqcup\{f_{ab} : b \sqsubseteq f(a) \ \wedge \ f_{ab} \text{ is strongly finite}\}) \sqsubseteq f$

  Suppose $b \sqsubseteq f(a)$. There are two cases.

  If $a \sqsubseteq x$, then $f_{ab}(x) = b \sqsubseteq f(a) \sqsubseteq f(x)$ by monotonicity of f.

  If $a \not\sqsubseteq x$, then $f_{ab}(x) = \bot \sqsubseteq f(x)$.

  Thus $f_{ab} \sqsubseteq f$.

- $f \sqsubseteq \bigsqcup\{f_{ab} : b \sqsubseteq f(a) \ \wedge \ f_{ab} \text{ is strongly finite}\}$

  Given any x, consider all strongly finite elements $a \sqsubseteq x$. Since D_1 is algebraic, x is the l.u.b. of such a's. Since f is continuous, $f(x)$ is the l.u.b. of the $f(a)$'s for such a's. Since D_2 is algebraic, each such $f(a)$ is the l.u.b. of all strongly finite elements $b \sqsubseteq f(a)$.

  Thus $f(x)$ is the l.u.b. of all strongly finite b's such that $b \sqsubseteq f(a)$, for some strongly finite a such that $a \sqsubseteq x$. Hence, $f(x)$ is the l.u.b. of all $f_{ab}(x)$, for all strongly finite f_{ab} such that $b \sqsubseteq f(a)$. □

Exercise 6.3.31 *Given an algebraic c.p.o. D_1 and a c.p.o. D_2, a function $f : D_1 \to D_2$ is continuous if and only it is determined by its behavior on the strongly finite elements, in the sense that*

$$f(x) = \bigsqcup\{f(a) : a \sqsubseteq x \ \wedge \ a \textit{ strongly finite}\}.$$

(Hint: one direction is immediate by continuity. Conversely, if a function is determined by its behavior on the strongly finite elements, then it is obviously monotone, and hence $\bigsqcup f(X) \sqsubseteq f(\bigsqcup X)$. Conversely, if a is strongly finite and $a \sqsubseteq \bigsqcup X$, then $a \sqsubseteq x$ for some $x \in X$, and so

$$\begin{aligned}
f(\bigsqcup X) &= \bigsqcup\{f(a) : a \sqsubseteq \bigsqcup X \ \wedge \ a \text{ strongly finite}\} \\
&\sqsubseteq \bigsqcup\{f(a) : a \sqsubseteq x \ \wedge \ a \text{ strongly finite} \ \wedge \ x \in X\} \\
&= \bigsqcup\{f(x) : x \in X\} \\
&= \bigsqcup f(X).
\end{aligned}$$

Thus f is continuous.)

Subcategories

representations (as for Heyting algebras)?

æ

Chapter 7

The Lawvere-Lambek Isomorphism

Implicit in the proof of the Categorical Soundness Theorem there is an actual translation of proofs in Natural Deductions into categorical language. In this section we show how to make this translation explicit, thus providing an isomorphism of languages.

7.1 Equational Presentation of Cartesian Closed Categories

In a way similar to what we did for Heyting $\sqcap$-algebras, we can present the theory of cartesian closed categories in a purely equational way.

Proposition 7.1.1 Equational Presentation of Cartesian Closed Categories (Lambek [1968], [1969], [1972]) *A categorical structure*

$$\mathcal{C} = \langle Ob, Hom, \times, \Rightarrow, 1 \rangle$$

is cartesian closed if and only if there are morphisms

$$id_A \quad t_A \quad l_{A,B} \quad r_{A,B} \quad eval_{A,B}$$

and functions

$$\circ \quad \langle\ \rangle_{C,A,B} \quad curry_{C,A,B}$$

of the appropriate types, i.e. satisfying the following **axioms***:*

- $id_A : A \to A$

- $t_A : A \to 1$
- $l_{A,B} : A \times B \to A$
- $r_{A,B} : A \times B \to B$
- $eval_{A,B} : (A \Rightarrow B) \times A \to B$

and the following **rules***:*

- $\dfrac{f : A \to B \quad g : B \to C}{g \circ f : A \to C}$
- $\dfrac{f : C \to A \quad g : C \to B}{\langle f, g\rangle_{C,A,B} : C \to A \times B}$
- $\dfrac{f : C \times A \to B}{curry_{C,A,B}(f) : C \to (A \Rightarrow B),}$

and such that the following **equations** *hold:*

1. $f = t_A$, *for all* $f : A \to 1$
2. $l_{A,B} \circ \langle f, g\rangle_{C,A,B} = f$
3. $r_{A,B} \circ \langle f, g\rangle_{C,A,B} = g$
4. $\langle l_{A,B} \circ h, r_{A,B} \circ h\rangle_{C,A,B} = h$
5. $eval_{A,B} \circ (curry_{C,A,B}(f) \times id_A) = f$
6. $curry_{C,A,B}(eval_{A,B} \circ (h \times id_A)) = h.$

Proof. The conditions are obviously necessary. To show that they are sufficient, we have to prove the following:

- 1 *is the terminal object*

  Property 1 says that there is a unique morphism from A to 1.

- $\times$ *is the categorical product*

  Properties 2 and 3 are the same as in the definition of product. It is thus enough to show that $\langle f, g\rangle$ is the unique morphism h such that

  $$f = l_{A,B} \circ h \qquad \text{and} \qquad g = r_{A,B} \circ h.$$

  By property 4,

  $$\langle f, g\rangle = \langle l_{A,B} \circ h, r_{A,B} \circ h\rangle = h.$$

- $\Rightarrow$ *is the categorical exponential*

  Property 5 is the same as in the definition of exponential. It is thus enough to show that *curry* (f) is the unique morphism h such that

  $$f = eval \circ (h \times id_A).$$

  By property 6,

  $$curry\ (f) = curry\ (eval \circ (h \times id_A)) = h. \quad \Box$$

7.2 Natural Deduction and Categories

We now examine more closely the inductive procedure in the proof of the Categorical Soundness Theorem, that produces a morphism $f : [\![\Gamma]\!] \to [\![\alpha]\!]$ associated to any derivation $\Gamma \vdash_{\mathcal{N}} \alpha$.

The language of Category Theory

Since we are dealing here not with a specific category, but only with the general language of category theory, we will not talk of interpreting formulas, but rather of translating them.

Thus, instead of associating to a formula α an object $[\![\alpha]\!]$, we will associate to it a letter A in the language of category theory, and use the convention that Latin letters are associated to corresponding Greek letters.

We will keep the correspondence between $\rightarrow$ and $\wedge$ on the one hand, and $\Rightarrow$ and $\times$ on the other. In particular, a proof of α from premises $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ will be translated into a morphism from $C = C_1 \times \cdots \times C_n$ to A.

Premises

The first part of the translation turns the assumptions of Natural Deduction into identities. This can be done in two ways, depending on which presentation of Natural Deduction we choose.

If we consider proof trees, then the axioms are simply formulas, and their translations will be identity functions. In particular, a premise α is translated as $id_A : A \rightarrow A$.

If we consider derivations instead, then the axioms take the form $\Gamma \vdash_{\mathcal{N}} \alpha$, where Γ is a set of formulas containing α. The translation involves here two steps: first we translate Γ as a cartesian product, and then we translate $\Gamma \vdash_{\mathcal{N}} \alpha$ as a projection from a product to one of its factors.

We have to argue first that the translation is possible, and then that it is well defined. The projections

$$p^{n,i}_{A_1,\dots,A_n} : A_1 \times \cdots \times A_n \to A_i$$

can be defined inductively, as follows:

$$\begin{array}{rcl} p^{1,1}_{A_1} & = & id_{A_1} \\ p^{n+1,1}_{A_1,\dots,A_{n+1}} & = & l_{A_1,A_2\times\cdots\times A_{n+1}} \\ p^{n+1,i+1}_{A_1,\dots,A_{n+1}} & = & p^{n,i}_{A_2,\dots,A_{n+1}} \circ r_{A_1,A_2\times\cdots\times A_{n+1}} \end{array}$$

where, obviously,

$$l_{A,B} : A \times B \to A \qquad \text{and} \qquad r_{A,B} : A \times B \to B.$$

These projections are well-defined because of 6.1.12.a and 6.1.12.b, asserting associativity and commutativity of the cartesian product.

Notice that the case $n = 1$ reduces to the case of a single assumption, and it is treated consistently with it, by using the appropriate identity function.

Introduction and elimination rules

The second part of the translation turns the introduction and elimination rules of Natural Deduction into axioms and rules of cartesian closed categories.

We first look at conjunction, and translate the left-hand-side into the right-hand-side:

$$\frac{\begin{array}{cc} \Gamma & \Gamma \\ \mathcal{D}_1 & \mathcal{D}_2 \\ \alpha & \beta \end{array}}{\alpha \wedge \beta} \qquad \frac{f_1 : C \to A \qquad f_2 : C \to B}{\langle f_1, f_2 \rangle : C \to A \times B}$$

$$\frac{\begin{array}{c} \Gamma \\ \mathcal{D} \\ \alpha \wedge \beta \end{array}}{\alpha} \qquad \frac{f : C \to A \times B}{l_{A,B} \circ f : C \to A}$$

$$\frac{\begin{array}{c} \Gamma \\ \mathcal{D} \\ \alpha \wedge \beta \end{array}}{\beta} \qquad \frac{f : C \to A \times B}{r_{A,B} \circ f : C \to B.}$$

We then look at implication, and translate the left-hand-side into the right-hand-side:

$$\dfrac{\begin{array}{c}\Gamma, [\alpha]^{(1)}\\ \mathcal{D}\\ \beta\end{array}}{\alpha^{(1)} \to \beta} \qquad\qquad \dfrac{f : C \times A \to B}{\mathit{curry}\ (f) : C \to (A \Rightarrow B)}$$

$$\dfrac{\begin{array}{cc}\Gamma & \Gamma\\ \mathcal{D}_1 & \mathcal{D}_2\\ \alpha \to \beta & \alpha\end{array}}{\beta} \qquad\qquad \dfrac{f_1 : C \to (A \Rightarrow B) \quad\quad f_2 : C \to A}{\mathit{eval}_{A,B} \circ \langle f_1, f_2\rangle : C \to B.}$$

An example

Before proceeding further, we give an example of how the previous translation is already sufficient to code proofs in Natural Deduction into the categorical language. By letting

$$\Gamma = \{\alpha \to (\beta \to \gamma), \alpha \to \beta, \alpha\},$$

we consider

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\Gamma \vdash \alpha \to (\beta \to \gamma) \quad\quad \Gamma \vdash \alpha}{\Gamma \vdash \beta \to \gamma} \qquad \dfrac{\Gamma \vdash \alpha \to \beta \quad\quad \Gamma \vdash \alpha}{\Gamma \vdash \beta}}{\alpha \to (\beta \to \gamma), \alpha \to \beta, \alpha \vdash \gamma}}{\alpha \to (\beta \to \gamma), \alpha \to \beta \vdash \alpha \to \gamma}}{\alpha \to (\beta \to \gamma) \vdash (\alpha \to \beta) \to (\alpha \to \gamma)}}{\vdash [\alpha \to (\beta \to \gamma)] \to [(\alpha \to \beta) \to (\alpha \to \gamma)].}$$

By letting

$$C = (A \Rightarrow (B \Rightarrow D)) \times (A \Rightarrow B) \times A,$$

this is translated into:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{p_C^{3,1} \quad\quad p_C^{3,3}}{\mathit{eval} \circ \langle p_C^{3,1}, p_C^{3,3}\rangle} \qquad \dfrac{p_C^{3,2} \quad\quad p_C^{3,3}}{\mathit{eval} \circ \langle p_C^{3,2}, p_C^{3,3}\rangle}}{\mathit{eval} \circ \langle \mathit{eval} \circ \langle p_C^{3,1}, p_C^{3,3}\rangle, \mathit{eval} \circ \langle p_C^{3,2}, p_C^{3,3}\rangle\rangle}}{\mathit{curry}\ (\mathit{eval} \circ \langle \mathit{eval} \circ \langle p_C^{3,1}, p_C^{3,3}\rangle, \mathit{eval} \circ \langle p_C^{3,2}, p_C^{3,3}\rangle\rangle)}}{\mathit{curry}\ (\mathit{curry}\ (\mathit{eval} \circ \langle \mathit{eval} \circ \langle p_C^{3,1}, p_C^{3,3}\rangle, \mathit{eval} \circ \langle p_C^{3,2}, p_C^{3,3}\rangle\rangle))}}{\mathit{curry}\ (\mathit{curry}\ (\mathit{curry}\ (\mathit{eval} \circ \langle \mathit{eval} \circ \langle p_C^{3,1}, p_C^{3,3}\rangle, \mathit{eval} \circ \langle p_C^{3,2}, p_C^{3,3}\rangle\rangle))).}$$

Normalization steps

The third part of the translation turns the normalization steps of Natural Deduction into equations of cartesian closed categories.

There are two normalization rules for conjunction, according to whether the left or right formula of a conjunction is eliminated.

The step

from

$$\dfrac{\dfrac{\begin{matrix}\Gamma & \Gamma\\ \mathcal{D}_1 & \mathcal{D}_2\\ \alpha & \beta\end{matrix}}{\alpha \wedge \beta}}{\alpha} \qquad\qquad \dfrac{\dfrac{f_1 : C \to A \quad f_2 : C \to B}{\langle f_1, f_2\rangle : C \to A \times B}}{l_{A,B} \circ \langle f_1, f_2\rangle : C \to A}$$

to

$$\begin{matrix}\Gamma\\ \mathcal{D}_1\\ \alpha\end{matrix} \qquad\qquad f_1 : C \to A$$

corresponds to equation 7.1.1.2:

$$l_{A,B} \circ \langle f_1, f_2\rangle = f_1.$$

Similarly, the step

from

$$\dfrac{\dfrac{\begin{matrix}\Gamma & \Gamma\\ \mathcal{D}_1 & \mathcal{D}_2\\ \alpha & \beta\end{matrix}}{\alpha \wedge \beta}}{\beta} \qquad\qquad \dfrac{\dfrac{f_1 : C \to A \quad f_2 : C \to B}{\langle f_1, f_2\rangle : C \to A \times B}}{r_{A,B} \circ \langle f_1, f_2\rangle : C \to B}$$

to

$$\begin{matrix}\Gamma\\ \mathcal{D}_2\\ \beta\end{matrix} \qquad\qquad f_2 : C \to B$$

corresponds to equation 7.1.1.3:

$$r_{A,B} \circ \langle f_1, f_2\rangle = f_2.$$

Finally, the step

from

$$\dfrac{\dfrac{\begin{array}{c}\Gamma,[\alpha]^{(1)}\\ \mathcal{D}_1\\ \beta\end{array}}{\alpha^{(1)}\to\beta} \qquad \begin{array}{c}\Gamma\\ \mathcal{D}_2\\ \alpha\end{array}}{\beta} \qquad\qquad \dfrac{\dfrac{f_1 : C\times A\to B}{\mathit{curry}\ (f_1) : C\to(A\Rightarrow B)} \qquad f_2 : C\to A}{\mathit{eval}\circ\langle \mathit{curry}\ (f_1), f_2\rangle : C\to B}$$

to

$$\begin{array}{c}\Gamma\\ \mathcal{D}_2\\ \Gamma,\quad \alpha\\ \mathcal{D}_1\\ \beta\end{array} \qquad\qquad f_1\circ\langle \mathit{id}_C, f_2\rangle : C\to B$$

corresponds to the following equation:

$$\mathit{eval}\circ\langle \mathit{curry}\ (f_1), f_2\rangle = f_1\circ\langle \mathit{id}_C, f_2\rangle.$$

Since this equation is not exactly 7.1.1.5, we now show that it actually follows from it.

Proposition 7.2.1 *The equation*

$$\mathit{eval}\circ\langle \mathit{curry}\ (f), g\rangle = f\circ\langle \mathit{id}, g\rangle$$

is a consequence of 7.1.1.5.

Proof. It is enough to notice that

$$\begin{array}{rl} & \mathit{eval}\circ\langle \mathit{curry}\ (f), g\rangle\\ = & \mathit{eval}\circ(\mathit{curry}\ (f)\times \mathit{id})\circ\langle \mathit{id}, g\rangle\\ = & f\circ\langle \mathit{id}, g\rangle \end{array}$$

by 6.1.15.d and 7.1.1.5. □

Equation 7.1.1.5

$$\mathit{eval}\circ(\mathit{curry}\ (f)\times \mathit{id}_A) = f$$

corresponds not to the general normalization step, but to the following special case of it:

from

$$\dfrac{\dfrac{\begin{matrix}\Gamma, [\alpha]^{(1)}\\ \mathcal{D}\\ \beta\end{matrix}}{\alpha^{(1)} \to \beta} \qquad \alpha}{\beta} \qquad\qquad \dfrac{\dfrac{f : C \times A \to B}{\mathit{curry}\ (f) : C \to (A \Rightarrow B)} \qquad id_A : A \to A}{\mathit{eval} \circ (\mathit{curry}\ (f) \times id_A) : C \times A \to B}$$

to

$$\begin{matrix}\Gamma, \alpha\\ \mathcal{D}\\ \beta\end{matrix} \qquad\qquad f : C \times A \to B.$$

Notice how, since here the two first subderivations do not have the same sets of premises, we are forced to use the operation $f \times g$ in place of $\langle f, g\rangle$. The same would hold in general, if our presentation of Natural Deduction allowed different sets of premises in the rules. For example, the $\wedge$-introduction rule would then become

$$\dfrac{\begin{matrix}\Gamma_1\\ \mathcal{D}_1\\ \alpha\end{matrix} \quad \begin{matrix}\Gamma_2\\ \mathcal{D}_2\\ \beta\end{matrix}}{\alpha \wedge \beta} \qquad \dfrac{f_1 : C_1 \to A \qquad f_2 : C_2 \to B}{f_1 \times f_2 : C_1 \times C_2 \to A \times B.}$$

Similarly for the →-introduction rule, and the two normalization rules.

Symmetric normalization steps

At this point we have translated all axioms and rules of Natural Deduction in the categorical language, but we have left out the crucial equations 7.1.1.4 and 7.1.1.6, which were introduced to ensure uniqueness of products and exponentials.

As the previous ones, these equations too corresponds to normalization steps in natural deduction, symmetric to the ones considered so far. They get rid not of *introductions followed by eliminations*, but rather of *eliminations followed by introductions*.

The symmetric normalization step for conjuction:

from

$$\dfrac{\dfrac{\begin{matrix}\Gamma\\ \mathcal{D}\\ \alpha \wedge \beta\end{matrix}}{\alpha} \qquad \dfrac{\begin{matrix}\Gamma\\ \mathcal{D}\\ \alpha \wedge \beta\end{matrix}}{\beta}}{\alpha \wedge \beta} \qquad\qquad \dfrac{\dfrac{h : C \rightarrow A \times B}{l_{A,B} \circ h : C \rightarrow A} \quad \dfrac{h : C \rightarrow A \times B}{r_{A,B} \circ h : C \rightarrow B}}{\langle l_{A,B} \circ h, r_{A,B} \circ h \rangle : C \rightarrow A \times B}$$

to

$$\begin{matrix}\Gamma\\ \mathcal{D}\\ \alpha \wedge \beta\end{matrix} \qquad\qquad h : C \rightarrow A \times B,$$

corresponds to equation 7.1.1.4:

$$\langle l_{A,B} \circ h, r_{A,B} \circ h \rangle = h.$$

The symmetric normalization step for implication:

from

$$\dfrac{\dfrac{\begin{matrix}\Gamma\\ \mathcal{D}\\ \alpha \rightarrow \beta\end{matrix} \qquad [\alpha]^{(1)}}{\beta}}{\alpha^{(1)} \rightarrow \beta} \qquad\qquad \dfrac{\dfrac{h : C \rightarrow (A \Rightarrow B) \quad id_A : A \rightarrow A}{eval \circ (h \times id_A) : C \times A \rightarrow B}}{curry(eval \circ (h \times id_A)) : C \rightarrow (A \Rightarrow B)}$$

to

$$\begin{matrix}\Gamma\\ \mathcal{D}\\ \alpha \rightarrow \beta\end{matrix} \qquad\qquad h : C \rightarrow (A \Rightarrow B),$$

corresponds to equation 7.1.1.6:

$$curry(eval \circ (h \times id_A)) = h.$$

The isomorphism

We can now finally state the result we have proved.

Theorem 7.2.2 The Lawvere-Lambek Isomorphism (Mann [1975], Seely [1977], [1983]) *There is an isomorphism between:*

- *the intuitionistic proof theory of implication and conjunction, with the Natural Deduction rules of introduction, elimination, normalization and inverse normalization;*
- *the equational theory of cartesian closed categories.*

7.3 Functional Completeness

The isomorphism between Natural Deduction and Cartesian Closed Categories immediately raises the question of the relationship of the latter with other formulations of logic, in particular Hilbert systems.

Theorem 7.3.1 Functional Completeness (Lambek [1972], [1974]) *For every indeterminate $x : 1 \to A$ and polynomial $\varphi(x) : 1 \to C$ over a cartesian closed category, it is possible to factor x out of $\varphi(x)$. More precisely, there are unique morphisms*

$$\begin{array}{lll} g : A \to C & \textit{such that} & \varphi(x) = g \circ x \\ h : 1 \to (A \Rightarrow C) & \textit{such that} & \varphi(x) = eval \circ \langle h, x \rangle. \end{array}$$

Proof. First we notice that it is enough to show that there exists a unique polynomial $\varphi^* : A \times 1 \to C$ not containing x, and such that

$$\varphi(x) = \varphi^* \circ \langle x, id_1 \rangle,$$

where

$$\begin{tikzcd} 1 \arrow[r, "\langle x, id_1 \rangle"] & A \times 1 \arrow[r, "\varphi^*"] & C. \end{tikzcd}$$

From φ^* we can easily obtain the following:

- *definition of g*

  By the properties of 1, the identity function id_1 can be decomposed as follows:

  $$\begin{tikzcd} 1 \arrow[r, "x"] & A \arrow[r, "t_A"] & 1. \end{tikzcd}$$

  Then

  $$\begin{array}{rcl} \varphi(x) & = & \varphi^* \circ \langle x, id_1 \rangle \\ & = & \varphi^* \circ \langle x, t_A \circ x \rangle \\ & = & \varphi^* \circ \langle id_A, t_A \rangle \circ x \end{array}$$

  by 6.1.12.d, and it is enough to let

  $$g = \varphi^* \circ \langle id_A, t_A \rangle.$$

- *definition of h*

  Having now
  $$\begin{tikzcd} 1 \times A \arrow[r, "r_{1,A}"] & A \arrow[r, "g"] & C, \end{tikzcd}$$
  we get
  $$\begin{tikzcd} 1 \arrow[r, "curry\ (g \circ r_{1,A})"] & A \Rightarrow C. \end{tikzcd}$$
  By letting
  $$h = curry\ (g \circ r_{1,A}),$$
  we have
  $$\begin{array}{rcl} eval \circ \langle h, x \rangle & = & eval \circ (h \times id) \circ \langle id, x \rangle \\ & = & eval \circ (curry\ (g \circ r_{1,A}) \times id) \circ \langle id, x \rangle \\ & = & g \circ r_{1,A} \circ \langle id, x \rangle \\ & = & g \circ x \\ & = & \varphi(x) \end{array}$$
  by 6.1.15.d.

We thus concentrate on the definition of φ^*, by induction on its construction. To make the induction possible we need to prove a slightly more general fact, namely: *for every indeterminate $x : 1 \to A$ and polynomial $\varphi(x) : B \to C$ over a cartesian closed category, there is a unique polynomial $\varphi^* : A \times B \to C$ not containing x, and such that*
$$\varphi(x) = \varphi^* \circ \langle x \circ t_B, id_B \rangle,$$
where
$$\begin{tikzcd} B \arrow[r, "\langle x \circ t_B, id_B \rangle"] & A \times B \arrow[r, "\varphi^*"] & C. \end{tikzcd}$$

The proof procedes in a number of steps.

1. *definition of φ^**

   By induction on the construction of $\varphi(x)$, we let:
   $$\varphi^* = \begin{cases} \varphi(x) \circ r_{A,B} & \text{if } \varphi(x) \text{ does not contain } x \\ l_{A,1} & \text{if } \varphi(x) = x \\ \varphi_1^* \circ \langle l_{A,B}, \varphi_2^* \rangle & \text{if } \varphi(x) = \varphi_1(x) \circ \varphi_2(x) \\ \langle \varphi_1^*, \varphi_2^* \rangle & \text{if } \varphi(x) = \langle \varphi_1(x), \varphi_2(x) \rangle \\ curry\ (\varphi_1^* \circ a_{A,B,C_1}) & \text{if } \varphi(x) = curry\ \varphi_1(x), \end{cases}$$
   where
   $$a_{A,B,C_1} : (A \times B) \times C_1 \longrightarrow A \times (B \times C_1)$$
   is the obvious isomorphism given by associativity of product (see 6.1.12.b).

   We first check that the definition makes sense, i.e. that $\varphi^* : A \times B \to C$.

- If $\varphi(x)$ does not contain x, then
$$A \times B \xrightarrow{r_{A,B}} B \xrightarrow{\varphi(x)} C.$$
- If $\varphi(x) = x$, then
$$A \times 1 \xrightarrow{l_{A,1}} A.$$
- If $\varphi(x) = \varphi_1(x) \circ \varphi_2(x)$ and
$$B \xrightarrow{\varphi_2(x)} D \xrightarrow{\varphi_1(x)} C,$$
then
$$A \times B \xrightarrow{\varphi_2^*} D \qquad \text{and} \qquad A \times D \xrightarrow{\varphi_1^*} C$$
by the induction hypothesis. So
$$A \times B \xrightarrow{\langle l_{A,B}, \varphi_2^* \rangle} A \times D \xrightarrow{\qquad \varphi_1^* \qquad} C.$$
- If $\varphi(x) = \langle \varphi_1(x), \varphi_2(x) \rangle$ and
$$B \xrightarrow{\varphi_1(x)} C_1 \qquad \text{and} \qquad B \xrightarrow{\varphi_2(x)} C_2,$$
then
$$A \times B \xrightarrow{\varphi_1^*} C_1 \qquad \text{and} \qquad A \times B \xrightarrow{\varphi_2^*} C_2$$
by the induction hypothesis. So
$$A \times B \xrightarrow{\langle \varphi_1^*, \varphi_2^* \rangle} C_1 \times C_2.$$
- If $\varphi(x) = \mathit{curry}\ \varphi_1(x)$ and
$$B \times C_1 \xrightarrow{\varphi_1(x)} C_2,$$
then
$$(A \times B) \times C_1 \xrightarrow{a_{A,B,C_1}} A \times (B \times C_1) \xrightarrow{\qquad \varphi_1^* \qquad} C_2$$
by the induction hypothesis. So
$$A \times B \xrightarrow{\mathit{curry}\ (\varphi_1^* \circ a_{A,B,C_1})} C_1 \Rightarrow C_2.$$

2. *existence of* φ^*

   We now check that φ^* defined above does indeed satisfy

   $$\varphi(x) = \varphi^* \circ \langle x \circ t_B, id_B\rangle.$$

   - If $\varphi(x)$ does not contain x, then

     $$\begin{aligned}\varphi^* \circ \langle x \circ t_B, id_B\rangle &= \varphi(x) \circ r_{A,B} \circ \langle x \circ t_B, id_B\rangle \\ &= \varphi(x) \circ id_B \\ &= \varphi(x).\end{aligned}$$

   - If $\varphi(x) = x$, then

     $$\begin{aligned}\varphi^* \circ \langle x \circ t_B, id_B\rangle &= l_{A,1} \circ \langle x \circ t_B, id_B\rangle \\ &= x \circ t_B \\ &= x \circ id_1 \\ &= x,\end{aligned}$$

     where $t_B = id_1$ because $B = 1$, and both t_1 and id_1 are morphisms from 1 to 1 (by 7.1.1.1, there can be only one such morphism).

   - If $\varphi(x) = \varphi_1(x) \circ \varphi_2(x)$, then

     $$\begin{aligned}\varphi^* \circ \langle x \circ t_B, id_B\rangle &= \varphi_1^* \circ \langle l_{A,B}, \varphi_2^*\rangle \circ \langle x \circ t_B, id_B\rangle \\ &= \varphi_1^* \circ \langle l_{A,B} \circ \langle x \circ t_B, id_B\rangle, \varphi_2^* \circ \langle x \circ t_B, id_B\rangle\rangle \\ &= \varphi_1^* \circ \langle x \circ t_B, \varphi_2(x)\rangle \\ &= \varphi_1^* \circ \langle x \circ t_D \circ \varphi_2(x), \varphi_2(x)\rangle \\ &= \varphi_1^* \circ \langle x \circ t_D, id_D\rangle \circ \varphi_2(x) \\ &= \varphi_1(x) \circ \varphi_2(x) \\ &= \varphi(x),\end{aligned}$$

     where everything follows by 6.1.12.d, 7.1.1.2, one induction hypothesis, the fact that $t_B = t_D \circ \varphi_2(x)$ because both are morphisms from B to 1, 6.1.12.d again, and the other induction hypothesis.

   - If $\varphi(x) = \langle \varphi_1(x), \varphi_2(x)\rangle$, then

     $$\begin{aligned}\varphi^* \circ \langle x \circ t_B, id_B\rangle &= \langle \varphi_1^*, \varphi_2^*\rangle \circ \langle x \circ t_B, id_B\rangle \\ &= \langle \varphi_1^* \circ \langle x \circ t_B, id_B\rangle, \varphi_2^* \circ \langle x \circ t_B, id_B\rangle\rangle \\ &= \langle \varphi_1(x), \varphi_2(x)\rangle \\ &= \varphi(x),\end{aligned}$$

     where everything follows by 6.1.12.d, and induction hypotheses.

- If $\varphi(x) = \mathit{curry}\ \varphi_1(x)$, then

$$\begin{array}{rl}
& \varphi^* \circ \langle x \circ t_B, id_B \rangle \\
= & \mathit{curry}\ (\varphi_1^* \circ a_{A,B,C_1}) \circ \langle x \circ t_B, id_B \rangle \\
= & \mathit{curry}\ (\varphi_1^* \circ a_{A,B,C_1} \circ (\langle x \circ t_B, id_B \rangle \times id_{C_1})) \\
= & \mathit{curry}\ (\varphi_1^* \circ a_{A,B,C_1} \circ (\langle \langle x \circ t_B, id_B \rangle \circ l_{B,C_1}, r_{B,C_1} \rangle)) \\
= & \mathit{curry}\ (\varphi_1^* \circ a_{A,B,C_1} \circ (\langle \langle x \circ t_B \circ l_{B,C_1}, l_{B,C_1} \rangle, r_{B,C_1} \rangle)) \\
= & \mathit{curry}\ (\varphi_1^* \circ \langle x \circ t_B \circ l_{B,C_1}, \langle l_{B,C_1}, r_{B,C_1} \rangle \rangle)) \\
= & \mathit{curry}\ (\varphi_1^* \circ \langle x \circ t_B \circ l_{B,C_1}, id_{B \times C_1} \rangle) \\
= & \mathit{curry}\ (\varphi_1^* \circ \langle x \circ t_{B \times C_1}, id_{B \times C_1} \rangle) \\
= & \mathit{curry}\ (\varphi_1(x)) \\
= & \varphi(x),
\end{array}$$

  where everything follows by 6.1.20.a, definition of ×, 6.1.12.d, definition of a (i.e. associativity of product), 6.1.12.e, the fact that $t_B \circ l_{B,C_1} = t_{B \times C_1}$ because both are morphisms from $B \times C_1$ to 1, and induction hypothesis.

3. *uniqueness of* φ^*

   We now show that if ψ is any polynomial not containing x and such that

$$\varphi(x) = \psi \circ \langle x \circ t_B, id_B \rangle,$$

   then $\psi = \varphi^*$.

   We first note that, by repeated applications of the clauses of the definition of φ^* indicated on the right:

$$\begin{array}{rll}
& (\psi \circ \langle x \circ t_B, id_B \rangle)^* & \\
= & \psi^* \circ \langle l_{A,B}, \langle x \circ t_B, id_B \rangle^* \rangle & \text{composition} \\
= & \psi \circ r_{A,B} \circ \langle l_{A,B}, \langle x \circ t_B, id_B \rangle^* \rangle & f \text{ (not containing } x) \\
= & \psi \circ \langle x \circ t_B, id_B \rangle^* & \\
= & \psi \circ \langle (x \circ t_B)^*, id_B^* \rangle & \langle\ \rangle \\
= & \psi \circ \langle x^* \circ \langle l_{A,B}, t_B^* \rangle, id_B^* \rangle & \text{composition} \\
= & \psi \circ \langle l_{A,1} \circ \langle l_{A,B}, t_B^* \rangle, id_B^* \rangle & x \\
= & \psi \circ \langle l_{A,B}, id_B^* \rangle & \\
= & \psi \circ \langle l_{A,B}, id_B \circ r_{A,B} \rangle & id_B \text{ (not containing } x) \\
= & \psi \circ \langle l_{A,B}, r_{A,B} \rangle & \\
= & \psi \circ id_{A \times B} & \text{by 6.1.15.a} \\
= & \psi. &
\end{array}$$

   This would show uniqueness, if we could claim at the beginning that

$$\varphi^* \quad = \quad (\psi \circ \langle x \circ t_B, id_B \rangle)^*.$$

Since we only supposed

$$\varphi(x) \quad = \quad \psi \circ \langle x \circ t_B, id_B \rangle,$$

we still need to prove that the $*$ operation preserves equality.

4. *preservation of equality*

   We proceed by induction on the definition of $=$, i.e. on the six clauses of 7.1.1, by repeated applications of the definition of $*$.

   - $f = t_A$, *for all* $f : A \to 1$

     Since $f^* = f \circ r_{A,B}$ and $t_A^* = t_A \circ r_{A,B}$ are both morphisms from $A \times B$ to 1, they must be equal by 7.1.1.1.

   - $l_{A,B} \circ \langle f, g \rangle = f$

$$\begin{array}{lll} & (l_{A,B} \circ \langle f, g \rangle)^* & \\ = & l_{A,B}^* \circ \langle l, \langle f, g \rangle^* \rangle & \\ = & l_{A,B} \circ r \circ \langle l, \langle f, g \rangle^* \rangle & \\ = & l_{A,B} \circ \langle f, g \rangle^* & \text{by } 7.1.1.3 \\ = & l_{A,B} \circ \langle f^*, g^* \rangle & \\ = & f^* & \text{by } 7.1.1.2. \end{array}$$

   - $r_{A,B} \circ \langle f, g \rangle = g$

     Similarly to the previous one.

   - $\langle l_{A,B} \circ h, r_{A,B} \circ h \rangle = h$

$$\begin{array}{lll} & \langle l_{A,B} \circ h, r_{A,B} \circ h \rangle^* & \\ = & \langle (l_{A,B} \circ h)^*, (r_{A,B} \circ h)^* \rangle & \\ = & \langle l_{A,B}^* \circ \langle l, h^* \rangle, r_{A,B}^* \circ \langle l, h^* \rangle \rangle & \\ = & \langle l_{A,B} \circ r \circ \langle l, h^* \rangle, r_{A,B} \circ r \circ \langle l, h^* \rangle \rangle & \\ = & \langle l_{A,B}^* \circ h^*, r_{A,B}^* \circ h^* \rangle & \text{by } 7.1.1.3 \\ = & h^* & \text{by } 7.1.1.4. \end{array}$$

   - $eval_{A,B} \circ (curry\ (f) \times id_A) = f$

$$\begin{array}{rll}
 & (\mathit{eval}_{A,B} \circ (\mathit{curry}\ (f) \times \mathit{id}_A))^* & \\
= & \mathit{eval}^*_{A,B} \circ \langle l, (\mathit{curry}\ (f) \times \mathit{id}_A)^* \rangle & \\
= & \mathit{eval}_{A,B} \circ r \circ \langle l, (\mathit{curry}\ (f) \times \mathit{id}_A)^* \rangle & \\
= & \mathit{eval}_{A,B} \circ (\mathit{curry}\ (f) \times \mathit{id}_A)^* & \text{by 7.1.1.3} \\
= & \mathit{eval}_{A,B} \circ \langle \mathit{curry}\ (f) \circ l, r \rangle^* & \text{by definition of } \times \\
= & \mathit{eval}_{A,B} \circ \langle (\mathit{curry}\ (f) \circ l)^*, r^* \rangle & \\
= & \mathit{eval}_{A,B} \circ \langle (\mathit{curry}\ f)^* \circ \langle l, l^* \rangle, r^* \rangle & \\
= & \mathit{eval}_{A,B} \circ \langle \mathit{curry}\ (f^* \circ a) \circ \langle l, l^* \rangle, r^* \rangle & \\
= & \mathit{eval}_{A,B} \circ (\mathit{curry}\ (f^* \circ a) \times \mathit{id}) \circ \langle \langle l, l^* \rangle, r^* \rangle & \text{by 6.1.15.c} \\
= & f^* \circ a \circ \langle \langle l, l^* \rangle, r^* \rangle & \text{by 7.1.1.5} \\
= & f^* \circ \langle l, \langle l^*, r^* \rangle \rangle & \text{by definition of } a \\
= & f^* \circ \langle l, \langle l, r \rangle^* \rangle & \\
= & f^* \circ \langle l, \mathit{id}^* \rangle & \text{by 6.1.12.e} \\
= & f^* \circ \langle l, \mathit{id} \circ r \rangle & \\
= & f^* \circ \langle l, r \rangle & \\
= & f^* \circ \mathit{id} & \text{by 6.1.12.e} \\
= & f^*. &
\end{array}$$

- $\mathit{curry}\ (\mathit{eval}_{A,B} \circ (h \times \mathit{id})) = h$

$$\begin{array}{rll}
 & \mathit{curry}^* (\mathit{eval}_{A,B} \circ (h \times \mathit{id})) & \\
= & \mathit{curry}\ ([\mathit{eval}_{A,B} \circ (h \times \mathit{id})]^* \circ a) & \\
= & \mathit{curry}\ (\mathit{eval}^*_{A,B} \circ \langle l, (h \times \mathit{id})^* \rangle \circ a) & \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ r \circ \langle l, (h \times \mathit{id})^* \rangle \circ a) & \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ (h \times \mathit{id})^* \circ a) & \text{by 7.1.1.3} \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ \langle h \circ l, r \rangle^* \circ a) & \text{by definition of } \times \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ \langle (h \circ l)^*, r^* \rangle \circ a) & \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ \langle (h^* \circ \langle l, l^* \rangle, r^* \rangle \circ a) & \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ (h^* \times \mathit{id}) \circ \langle \langle l, l^* \rangle, r^* \rangle \circ a) & \text{by 6.1.15.c} \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ (h^* \times \mathit{id}) \circ \langle l, \langle l^*, r^* \rangle \rangle) & \text{by definition of } a \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ (h^* \times \mathit{id}) \circ \mathit{id}) & \text{as above} \\
= & \mathit{curry}\ (\mathit{eval}_{A,B} \circ (h^* \times \mathit{id})) & \\
= & h^* & \text{by 7.1.1.6.} \quad \Box
\end{array}$$

Exercise 7.3.2 Generalized Functional Completeness. *For every indeterminate $x : D \to A$ and polynomial $\varphi(x) : D \to C$ over a cartesian closed category, there is a unique morphism $g : A \to C$ such that $\varphi(x) = g \circ x$.* (Hint: for every indeterminate $x : D \to A$ and polynomial $\varphi(x) : B \to C$, there is a unique polynomial $\varphi^* : (D \Rightarrow A) \times B \to C$ not containing x such that

$$\varphi(x) = \varphi^* \circ \langle \mathit{curry}\ (x \circ r_{B,D}), \mathit{id}_B \rangle,$$

where

$$\begin{tikzcd} B \arrow[r, "\langle curry\ (x \circ r_{B,D}), id_B \rangle"] & (D \Rightarrow A) \times B \arrow[r, "\varphi^*"] & C. \end{tikzcd}$$

The crucial difference with the case above is the definition of φ^* for the atomic cases, which is now:

$$\varphi^* = \begin{cases} \varphi(x) \circ r_{D \Rightarrow A,B} & \text{if } \varphi(x) \text{ does not contain } x \\ eval_{A,B} & \text{if } \varphi(x) = x.) \end{cases}$$

7.4 Symmetric Normalization

The isomorphism proved above does not allow us to deduce normalization and the Church-Rosser properties for the categorical equations, since we did not consider the symmetric normalization rules in the proof of the Normalization Theorem. We thus have to strengthen the result.

WARNING: church-rosser not yet discussed!

Characterization of models of typed lambda calculus as cartesian closed categories?

æ

Part C

Typed Lambda Calculus

Chapter 8

Syntax

In the previous chapters we have developed two substantially different but equivalent formalisms, namely the Intuitionistic Implicational Calculus with Conjunction and the theory of cartesian closed categories. We turn now to a third presentation of the same topic in terms of the Typed Lambda Calculus, which is a theory of functions with the following characteristics:

1. Functions are defined *intensionally*, by sets of rules that allow the computation of their values for given arguments.

2. Together with the definition of a function, a description is given of the *type* of both arguments and values.

The intensional approach is in accord both with classical (pre-Dirichlet) mathematical practice, as well as with the current needs of Computer Science, whose programs are intensional. But it is opposed to the modern, set-theoretical mathematical practice that defines a function f *extensionally* by means of its graph, i.e.

$$G_f = \{(x, y) : f(x) = y\}.$$

The specification of types for arguments and values is a logical consequence of the shift from the extensional back to the intensional approach, since in the former the graph automatically describes all possible arguments and values, and hence it implicitly defines domain and range, while in the latter an isolated rule defines in general a class of functions, rather than a single one. For example, the rule

'give as output the input itself'

computes the identity function on A, for any given set A.

8.1 Typed Lambda Terms

Instead of deriving the presentation of the Typed Lambda Calculus from the Intuitionistic Implicational Calculus, as we did for Heyting algebras and cartesian closed categories, we develop it independently in a parallel way. At the cost of some repetition, this provides an autonomous treatment that can be read independently of the previous chapters.

Types

Intuitively, a type is a set that can serve as the domain or the range of a function.

Atomic types represent atomic domains and ranges, i.e. sets whose elements are not further analyzed. As a first approximation to a theory of functions we will consider only unspecified atomic types, and denote them by type *letters*. In a second approximation we can also consider specific atomic types, representing particular sets of interest, and denote them by type *constants*, for example **Nat** for the natural numbers, **Real** for the real numbers, **Bool** for the Boolean algebra of truth values (variously denoted by $\{0, 1\}$, $\{F, T\}$ or $\{\bot, \top\}$). The presence of type constants distinguishes an *applied* theory of functions from a *pure* one.

Given types α and β, we can consider functions with arguments of type α and values of type β. These functions are objects of a different type, denoted by $\alpha \to \beta$.

The **language** for the description of types consists of:

- type letters $p, q, r, \ldots$
- parentheses '(' and ')'
- the type constructor $\to$ (*arrow*).

Types are defined inductively as follows:

- type letters are types (*atomic types*)
- if α and β are types, so is $(\alpha \to \beta)$.

To increase readability some parentheses can either be omitted, when no confusion arises, or written differently, e.g. as '[' and ']'. We will use lowercase Roman letters such as p for atomic types, and lowercase Greek letters such as α for types.

Terms

The intuitive way of looking at terms in a given language is to see them as *descriptions of objects*: more precisely, as *proper names* of individual objects, or as *generic names* of objects in a given set. For example, in Arithmetic a polynomial expression is a typical term, and it describes either a number (if it has no variables, as in $3^2 + 5 \cdot 7$) or a generic element in the range of an n-ary function (if it has n distinct variables, for example two in $x^2 + 5 \cdot y$).

In the Typed Lambda Calculus we are concerned with names needed in a theory of functions: names for the functions themselves, as well as names for their arguments and values.

To *define* a (unary) function as a rule of computation, we follow the usual mathematical practice: we exhibit a generic description of the value u^β of a given type β, describing the range, depending on a generic description of the argument x^α of a given type α, describing the domain. A subtle ambiguity occurs here: a description u^β in which x^α occurs can be taken either as a description of the value of a function for the generic argument x^α, or as a description of the function itself. It would be impossible to distinguish between the two uses, without a further specification of what is meant. For example, if x is a variable ranging over natural numbers, then $2x$ describes both a generic even *number*, and the *function* that associates to any number its double. These ambiguities are not precisely what we need in a rigorous theory, and some device is needed to distinguish between these two uses. We choose to always use a term u^β to denote an object of type β. When we want to write a name for the function

$$x^\alpha \longmapsto u^\beta$$

we will use the special symbol λ, and denote the function by

$$(\lambda x^\alpha .\, u^\beta)^{\alpha\to\beta}.$$

The subtle ambiguity has now been eliminated, but it is replaced by a subtle distinction: when we consider u^β as the description of a value, an occurrence of x^α in it is seen as the description of a generic argument; but after we have named the function a change in the status of x^α has happened, since $(\lambda x^\alpha .\, u^\beta)^{\alpha\to\beta}$ already contains the information that x^α is used as an argument. Technically, we say that x^α was (possibly) *free* in u^β and has become *bound* in $(\lambda x^\alpha .\, u^\beta)^{\alpha\to\beta}$. Only free variables are really variables, in the sense of describing generic objects, while bound variables are simply devices used to describe a function: they are used *in* a description, but are not *a* description.[1]

We not only want to be able to define functions, but also to *use* them. This is done by applying a function to an argument. Since a function $u^{\alpha\to\beta}$ has specified domain and range, the argument should be the description of an object of the right type. Thus $(u^{\alpha\to\beta} v^\alpha)^\beta$ will be a description of an object in the range of $u^{\alpha\to\beta}$, and hence of type β.

[1] In current mathematics, where the λ-notation is not used, we often drop the variable in the name of the value of a function to obtain the name of function itself, e.g. by stepping from '$\sin x$' to 'sin'.

We now proceed to give a formal inductive definition of typed λ-terms and, simultaneously, of their free and bound variables.

The **language** for the description of terms consists of:

- for any type α, *variables* $x^\alpha, y^\alpha, \ldots$
- parentheses '(' and ')'
- dot '.'
- the term constructor λ (*lambda operator*).

This language is enough for a first approximation to a theory of functions, in which we consider only unspecified atomic terms, and denote them by term *variables*. In a second approximation we can also consider specific atomic terms, representing particular objects or functions of interest, and denote them by term *constants*, for example an object **0** of atomic constant type **Nat**, or a function **Succ** of type **Nat** → **Nat**. The presence of term constants distinguishes an *applied* theory of functions from a *pure* one.

Definition 8.1.1 Typed λ-Terms (Church [1941]) *Typed λ-terms are defined inductively as follows.*

1. **Variables**. *A variable x^α is a term of type α, and x^α occurs free in it.*
2. **Functional Application**. *If $u^{\alpha\to\beta}$ and v^α are terms of type $\alpha \to \beta$ and α, then $(u^{\alpha\to\beta}v^\alpha)^\beta$ is a term of type β. An occurrence of a variable is free or bound in it if it was so in $u^{\alpha\to\beta}$ or v^α.*
3. **Functional Abstraction**. *If x^α is a variable of type α and u^β is a term of type β, then $(\lambda x^\alpha .\, u^\beta)^{\alpha\to\beta}$ is a term of type $\alpha \to \beta$. An occurrence of a variable is free or bound in it if it was so in u^β, with the exception of the free occurrences of x^α in u^β, which become bound.*

Terms in which no variable occurs free are called **closed**.

We should stress the fact that it is not *variables* that are free or bound in a term, but *occurrences* of them. In particular, a variable may occur both free and bound in the same term. For example, in $(\lambda x^\alpha .\, x^\alpha)^{\alpha\to\alpha}x^\alpha$ the last occurrence of x^α is free, the first one is bound.

To increase readability some parentheses can either be omitted, when no confusion arises, or written differently, e.g. as '[' and ']'. Similarly for the explicit indications of types. We will use the letters x, y, z, … for variables, and t, u, v, … for terms.

Functions of many variables

The λ-operator allows us to abstract only on a single variable, and thus to define only unary functions (of one argument). But any self-respecting theory of functions should be able to deal with functions of many arguments as well, since they often arise in common practice.

The main observation here is that *the presence of all type levels allows us to replace functions of many variables by iterated unary functions*. Even if at first sight this may sound strange, it is used in familiar practice. For example, in the classical inductive definition of the sum S of integers, due to Grassmann [1861]:

$$\begin{cases} S(x,0) & = & x \\ S(x,y+1) & = & S(x,y)+1. \end{cases}$$

This defines the *binary* function $S(x,y)$ by induction on only *one* variable y, with x appearing in the definition as a parameter. In other words, for any fixed x we are actually defining the *unary* function

$$S_x : y \longmapsto S(x,y)$$

as follows:

$$\begin{cases} S_x(0) & = & x \\ S_x(y+1) & = & S_x(y)+1. \end{cases}$$

The definition of the binary function S is thus reduced to the definition of a family $\{S_x\}_{x\in N}$ of unary functions. But a family of unary functions indexed by numbers is simply a higher type function from numbers to unary functions. Thus the binary function

$$(x,y) \longmapsto S(x,y)$$

from numbers to numbers is reduced to a unary function

$$x \longmapsto S_x \qquad \text{or} \qquad x \longmapsto (y \longmapsto S(x,y))$$

from numbers to unary functions, each of them from numbers to numbers.

Since in the Typed Lambda Calculus all type levels are available, we can play the same trick and identify functions of many variables with successive applications of unary functions, i.e. step from

$$(x_1^{\alpha_1}, x_2^{\alpha_2}, \ldots, x_n^{\alpha_n}) \longmapsto u^\beta$$

to

$$x_1^{\alpha_1} \longmapsto (x_2^{\alpha_2} \longmapsto \cdots (x_n^{\alpha_n} \longmapsto u^\beta)\cdots).$$

In our present notation, this corresponds to the following **convention on multiple $\boldsymbol{\lambda}$-abstractions**, that defines λ-abstraction on n-tuples of variables:

$$\lambda x_1^{\alpha_1} \cdots x_n^{\alpha_n} . u^\beta \overset{\text{def}}{=} \lambda x_1^{\alpha_1} . (\cdots (\lambda x_n^{\alpha_n} . u^\beta)\cdots),$$

where the right-hand-side has type

$$\alpha_1 \to (\cdots(\alpha_n \to \beta)\cdots).$$

A complementary **convention on multiple applications**, consistent with the previous one, allows us to consider the simultaneous application of a single term $t^{\alpha_1 \to (\cdots(\alpha_n \to \beta)\cdots)}$ to an n-tuple of terms with the appropriate types:

$$t^{\alpha_1 \to (\cdots(\alpha_n \to \beta)\cdots)} v_1^{\alpha_1} \cdots v_n^{\alpha_n} \overset{\text{def}}{=} (\cdots(tv_1)\cdots v_n),$$

where the right-hand-side has type β.

Since multiple abstractions are defined in terms of single abstractions, it makes sense to apply a term defined by an n-ary abstraction to less than n arguments: this only means that we perform less than n iterated applications.

Bound variables

As explained above, while a free occurence of a variable describes an object, although a generic one, a bound occurrence of a variable is simply a device indicating the argument of a function in the description of its generic value. It can thus be thought of as a pointer to an empty place waiting to be filled up by an argument, and the information it conveys should be independent of its name.

On the one hand, all variables of type α indicate a generic object of type α, and in this respect they are interchangeable. For example, the description $\lambda x^\alpha.\, x^\alpha$ of the identity function on objects of type α can be rephrased without any reference to x^α as 'the function which associates any object of type α to itself'. In particular, the same function would still be described by $\lambda y^\alpha.\, y^\alpha$, where y^α is any other variable of type α.

On the other hand, there is a difference in considering *two* generic objects of the same type, that need not be equal, and the same generic object *twice*. This is reflected in the distinction between occurrences of different variables of the same type, and different occurrences of the same variable. In other words, variables are not *completely* interchangeable. For example, the description $\lambda x^\alpha.\, y^\alpha$ can be rephrased as 'the constant function that associates to any object of type α the generic object y^α'. This reformulation does not contain any reference to x^α, but it still refers to y^α. In particular, the same function would still be described by $\lambda z^\alpha.\, y^\alpha$, but not by $\lambda y^\alpha.\, y^\alpha$: the latter describes the *identity* function on objects of type α, and not a *constant* function of unspecified value. The lesson of this example can be rephrased in the following general *credo*, to be followed while renaming bound variables:

whatever was free, should remain free.

The previous example can be pushed a bit further, by quantifying over the missing variable. If we consider $\lambda y^\alpha.\, (\lambda x^\alpha.\, y^\alpha)$, then we are describing the *one-one* function that to every object of type α associates the constant function with that element as value. But $\lambda x^\alpha.\, (\lambda x^\alpha.\, x^\alpha)$, in which the first bound variable has been renamed, describes now the *constant* function that to every element of type α associates the identity function. What happened is that y^α was free in the subterm $\lambda x^\alpha.\, y^\alpha$, but has become bound after the renaming. Thus the *credo* has to be interpreted in a strong form, as referring not only to global freedom in a term, but also to local freedom in all subterms.

We can now formulate a rule for correct renaming of bound variables. In practice, the easiest way to fulfill the conditions is to rename bound variables by variables that do not occur *at all* in a given term, either free or bound.

Definition 8.1.2 ***α*-Rule**. *In a given term we can change every bound occurrence of a variable with occurrences of another variable of the same type, as long as no free occurrence of any variable in any subterm of the original term becomes bound in that subterm after the change.*

The α-rule will be tacitly applied when needed, and we will not keep track of it. In other words, we will identify terms that can be obtained one from the other by correct applications of the α-rule, and consider them as inessential variations one of the others. Technically, this means that we are really considering as terms the equivalence classes of our original terms, under the equivalence relation induced by the α-rule.

Reductions

In Arithmetic, a polynomial expression can be evaluated at a given argument x, by plugging in some number for the variable. For example, $x^2 + 5 \cdot x$ can be evaluated at 3, with the result $3^2 + 5 \cdot 3$. More generally, we could evaluate the original polynomial expression at an argument that is still itself a polynomial expression, e.g. at y^3, thus getting $(y^3)^2 + 5 \cdot y^3$. Formally, these evaluation steps consist in the substitution of a description of a number for a variable.

The following rule allows us to do the same for the abstract terms introduced above. Intuitively it says that, since $\lambda x^\alpha.\, u^\beta$ denotes the function whose generic value is described by u^β (in terms of the generic description x^α of the variable), a description of the value of the function for a particular argument v^α can be obtained from the generic description of the value, by substituting the description v^α of a specific argument for the generic description of the argument (i.e., the variable x^α).

Definition 8.1.3 ***β*-Rule**. *Given terms u^β and v^α, we can step from $(\lambda x^\alpha.\, u^\beta)^{\alpha\to\beta} v^\alpha$ (called a* **redex***) to $u^\beta[x^\alpha := v^\alpha]$ (called a* **reduct**[2]*), where the latter is the result*

[2] Some authors call it a **contractum**.

of the substitution of v^α for the free occurrences of x^α in u^β. We write

$$(\lambda x^\alpha .\, u^\beta)v^\alpha \ \longrightarrow_{1\beta}\ u^\beta[x^\alpha := v^\alpha]$$

to state that one step of the β-rule has been applied to the left-hand-side to produce the right-hand-side.

Formally, $u^\beta[x^\alpha := v^\alpha]$ is defined by induction on u^β, as follows:

$$u^\beta[x^\alpha := v^\alpha] = \begin{cases} v^\alpha & \text{if } u^\beta = x^\alpha \\ u^\beta & \text{if } u^\beta = y^\beta \neq x^\alpha \\ (u_1^{\gamma\to\beta}[x^\alpha := v^\alpha])(u_2^\gamma[x^\alpha := v^\alpha]) & \text{if } u^\beta = u_1^{\gamma\to\beta}u_2^\gamma \\ \lambda y^\gamma .\, (u_1^\delta[x^\alpha := v^\alpha]) & \text{if } u^\beta = \lambda y^\gamma .\, u_1^\delta, \end{cases}$$

where in the last clause we tacitly use the α-rule to ensure that the bound variable is not x^α itself. Notice that, by induction, $u^\beta[x^\alpha := v^\alpha]$ has still type β.

The **α-rule** can be expressed, in terms of substitution, as follows (with the appropriate restrictions on y^α):

$$\lambda x^\alpha .\, u^\beta = \lambda y^\alpha .\, (u^\beta[x^\alpha := y^\alpha]).$$

Normal forms

Particularly simple polynomial expressions are the ones in *normal form*, i.e. the ones inside which no arithmetical reduction (using the definitions of sum and product) is possible:[3] in particular, the ones without variables are simply the decimal representations of numbers. The following well-known arithmetical fact is the fundamental computational property of simplification that will concern us here: *given a polynomial expression, any sequence of reductions will eventually produce one in normal form, which depends on the original expression but is independent of the chosen sequence of reductions*. The independence is useful in practice: for example, in a polynomial expression of the form $p_1 \cdot p_2$ such that one of p_1 and p_2 reduces to 0, there is an advantage in reducing it first, since then the whole expression would reduce to 0 in one more step. This requires freedom of choice in the reduction steps, and it would not be possible if we only knew that a *fixed* reduction procedure produces the normal form.

Sections 3 and 4 are devoted to showing that similar results hold for the Typed Lambda Calculus as well, in the following sense. A λ-term is said to be in **β-normal form** if no application of the β-rule is possible inside it, i.e. the term

[3] This obviously requires seeing the definitions of both sum and product as expressing not equalities, but reductions rules. For example,

$$(x + y) + 1 = x + (y + 1)$$

should be read as a reduction rule of the left-hand-side to the right-hand-side, but not conversely.

does not contain any redex. Then, *given a λ-term, any sequence of applications of the β-rule to it will eventually produce a term in β-normal form* (**Strong $\boldsymbol{\beta}$-Normalization Theorem**, 8.3.2 or 8.3.8), *that depends on the original term but is independent of the chosen sequence of β-reductions* (**Uniqueness of $\boldsymbol{\beta}$-Normal Forms**, 8.4.5 or 8.5.7).

The expression 'to apply the β-rule *inside* a term' should be intuitively clear. Formally, we need to extend the β-rule (as defined in 8.1.3) to allow for its application not only to a term that *is* a redex, but also to any term that *contains* a redex. For example, we should be allowed to step from

$$\lambda y^\gamma.\,[(\lambda x^{\alpha\to\beta}.\,x^{\alpha\to\beta})u^{\alpha\to\beta}]v^\alpha$$

to

$$\lambda y^\gamma.\,u^{\alpha\to\beta}v^\alpha.$$

The simplest way of doing this is by introducing rules that allow the reconstruction of a term from a reduct, in the same way the original term was constructed from a redex. Formally, we inductively extend the meaning of $\longrightarrow_{1\beta}$ as follows:

Definition 8.1.4 One-Step $\boldsymbol{\beta}$-Reducibility. *The reducibility $\longrightarrow_{1\beta}$ is defined inductively by the following clauses:*

$$(\lambda x^\alpha.\,u^\beta)v^\alpha \ \longrightarrow_{1\beta}\ u^\beta[x^\alpha := v^\alpha] \tag{8.1}$$

$$\frac{u_1^{\alpha\to\beta} \ \longrightarrow_{1\beta}\ u_2^{\alpha\to\beta}}{u_1^{\alpha\to\beta}v^\alpha \ \longrightarrow_{1\beta}\ u_2^{\alpha\to\beta}v^\alpha} \tag{8.2}$$

$$\frac{v_1^\alpha \ \longrightarrow_{1\beta}\ v_2^\alpha}{u^{\alpha\to\beta}v_1^\alpha \ \longrightarrow_{1\beta}\ u^{\alpha\to\beta}v_2^\alpha} \tag{8.3}$$

$$\frac{u_1^\beta \ \longrightarrow_{1\beta}\ u_2^\beta}{\lambda x^\alpha.\,u_1^\beta \ \longrightarrow_{1\beta}\ \lambda x^\alpha.\,u_2^\beta,} \tag{8.4}$$

where the first clause (i.e. the β-rule) can be thought of as an axiom, and the remaining ones as deduction rules.

Now the previous example can be dealt with by successively applying rules 1, 2 and 4, as follows:

$$\cfrac{\cfrac{(\lambda x^{\alpha\to\beta}.\,x^{\alpha\to\beta})u^{\alpha\to\beta} \ \longrightarrow_{1\beta}\ u^{\alpha\to\beta}}{[(\lambda x^{\alpha\to\beta}.\,x^{\alpha\to\beta})u^{\alpha\to\beta}]v^\alpha \ \longrightarrow_{1\beta}\ u^{\alpha\to\beta}v^\alpha}}{\lambda y^\gamma.\,[(\lambda x^{\alpha\to\beta}.\,x^{\alpha\to\beta})u^{\alpha\to\beta}]v^\alpha \ \longrightarrow_{1\beta}\ \lambda y^\gamma.\,(u^{\alpha\to\beta}v^\alpha).}$$

A further extension, needed in the informal discussion above on normal forms, requires the application of the β-rule inside a given term not only once, but any finite number of times. This defines the notion of **$\boldsymbol{\beta}$-reducibility**, which we will indicate by $\longrightarrow_{\beta}$, and of which $\longrightarrow_{1\beta}$ constitutes a single step. By definition, $\longrightarrow_{\beta}$ is simply the *reflexive and transitive closure of* $\longrightarrow_{1\beta}$. Formally, it is defined as follows:

Definition 8.1.5 $\boldsymbol{\beta}$-Reducibility. *The reducibility $\longrightarrow_{\beta}$ is defined inductively by the following clauses:*

$$u^{\alpha} \longrightarrow_{\beta} u^{\alpha} \tag{8.5}$$

$$\frac{u_1^{\alpha} \longrightarrow_{1\beta} u_2^{\alpha}}{u_1^{\alpha} \longrightarrow_{\beta} u_2^{\alpha}} \tag{8.6}$$

$$\frac{u_1^{\alpha} \longrightarrow_{\beta} u_2^{\alpha} \qquad u_2^{\alpha} \longrightarrow_{\beta} u_3^{\alpha}}{u_1^{\alpha} \longrightarrow_{\beta} u_3^{\alpha},} \tag{8.7}$$

where the first clause can be thought of as an axiom, and the remaining ones as deduction rules.

To increase readability the subscripts 1β or β can be omitted, when no confusion arises.

Equality

The identification of normal forms with the objects dealt with in the Typed Lambda Calculus induces an identification of terms that have the same normal form, as descriptions of the same object. This could be taken as a definition of **$\boldsymbol{\beta}$-equality** for terms. Since normal forms always exist in the Typed Lambda Calculus, two terms have the same normal form if and only if they reduce to a common term. This could be taken as an equivalent definition of β-equality, working also for the Untyped Lambda Calculus (where normal forms do not always exist).

However, these definitions would not make it immediate to prove the fundamental property of equality, of being an equivalence relation. We thus turn things upside down and define β-equality, which we will indicate as $=_{\beta}$, as the equivalence relation generated by $\longrightarrow_{\beta}$, i.e. as the *symmetric and transitive closure of* $\longrightarrow_{\beta}$. Formally, $=_{\beta}$ is defined as follows:

Definition 8.1.6 $\boldsymbol{\beta}$-Equality. *The equivalence relation $=_{\beta}$ is defined inductively by the following clauses:*

$$\frac{u_1 \longrightarrow_{\beta} u_2}{u_1 =_{\beta} u_2} \tag{8.8}$$

$$\frac{u_1 =_{\beta} u_2}{u_2 =_{\beta} u_1} \tag{8.9}$$

$$\frac{u_1 =_{\beta} u_2 \qquad u_2 =_{\beta} u_3}{u_1 =_{\beta} u_3,} \tag{8.10}$$

where the first clause can be thought of as an axiom, and the remaining ones as deduction rules.

To increase readability, the subscript β can be omitted, when no confusion arises.

Proposition 8.1.7 *If two terms reduce to a common term, then they are β-equal.*

Proof. Suppose u_1 and u_2 both reduce to t. Then

$$\dfrac{\dfrac{u_1 \longrightarrow_\beta t}{u_1 =_\beta t} \qquad \dfrac{\dfrac{u_2 \longrightarrow_\beta t}{u_2 =_\beta t}}{t =_\beta u_2}}{u_1 =_\beta u_2}$$

is a derivation of $u_1 =_\beta u_2$ from rules 8.8 (twice), 8.9 and 8.10. □

Corollary 8.1.8 *If two terms have the same normal form, then they are β-equal.*

The converse of Corollary 8.1.8 follows from Existence and Uniqueness of Normal Forms, with a proof working only for the Typed Lambda Calculus (see 8.4.6). The converse of Proposition 8.1.7 follows from the Diamond Property, with a proof working also for the Untyped Lambda Calculus (see 8.5.8). Thus the proposed definitions of β-equality are all equivalent.

8.2 Combinators $\star$

Our first approach to terms for the description of functions has been *analytical* (top-down): we introduced an operator λ that allows us to name a unary function whenever we specify its argument, as well as a description of its generic value in terms of that argument.

We look now for a *synthetical* (bottom-up) approach, in which a few λ-terms called **atomic combinators** are selected, and **combinators** are built up from them by means of application alone. The question arises of finding nontrivial atomic combinators, such that the terms built up from them and the variables by means of application alone, which we will call **combinatorial terms**,[4] somehow represent all λ-terms. Since we do not know beforehand whether this is possible, and even if we did know it we would have no clue as to which atomic combinators would do the job, we just attempt a proof of the possibility. In the process of proving the result we will *discover* which atomic combinators are needed for the proof to go through. It is thus advisable for the reader to disregard on a first reading the statement of the next result, which would look completely unmotivated, and to come back to it only after having read the proof.

[4] The usual combinators are *closed* combinatorial terms, i.e. the ones without free variables.

Theorem 8.2.1 Combinatorial Completeness (Schönfinkel [1924], Curry [1930]) *Define the* **atomic combinators** *as follows, for any types α, β, and γ:*

- $\mathbf{I}^{\alpha\to\alpha} = \lambda x^{\alpha}.\, x^{\alpha}$
- $\mathbf{K}^{\gamma\to(\alpha\to\gamma)} = \lambda y^{\gamma} x^{\alpha}.\, y^{\gamma}$
- $\mathbf{S}^{[\alpha\to(\gamma\to\delta)]\to[(\alpha\to\gamma)\to(\alpha\to\delta)]} = \lambda x^{\alpha\to(\gamma\to\delta)} y^{\alpha\to\gamma} z^{\alpha}.\, (xz)^{\gamma\to\delta}(yz)^{\gamma}.$

Then for every λ-term t there is a **combinatorial term** t_c *built up from atomic combinators and variables by application alone, such that t_c is β-reducible to t.*

Proof. By definition 8.1.1, a λ-term t is built up from variables by application and λ-abstraction. If we want to get rid of the latter, we have to find special cases of it (the atomic combinators) that, together with the variables, are sufficient to generate all λ-abstractions by application.

We proceed inductively on the definition of λ-abstraction, and show how to turn $\lambda x^{\alpha}.\, u^{\beta}$ into a combinatorial term, when u^{β} is already such. The idea is to transform a definition of u^{β} as a combinatorial term into one of $\lambda x^{\alpha}.\, u^{\beta}$ by sticking 'λx^{α}.' in front of every subterm used in the original definition of u^{β}. Since u^{β} is built up from variables and atomic combinators by application, we have four cases to consider.

1. The occurrences of the variable x^{α} in the definition of u^{β} become occurrences of $\lambda x^{\alpha}.\, x^{\alpha}$, which is the atomic combinator $\mathbf{I}^{\alpha\to\alpha}$.

2. The occurrences of a variable y^{γ} different from x^{α} in the definition of u^{β} become occurrences of the term $(\lambda x^{\alpha}.\, y^{\gamma})^{\alpha\to\gamma}$. The idea is to postulate an atomic combinator $\mathbf{K}$ that would reduce to it when applied to the variable y^{γ}. We thus need a name for the function

$$y^{\gamma} \longmapsto \lambda x^{\alpha}.y^{\gamma}$$

and, using the convention of multiple λ-abstractions, this is precisely how $\mathbf{K}^{\gamma\to(\alpha\to\gamma)}$ is defined in 2 above. Then

$$\mathbf{K}^{\gamma\to(\alpha\to\gamma)} y^{\gamma} \;\longrightarrow_{\beta}\; \lambda x^{\alpha}.\, y^{\gamma}.$$

3. The occurrences of an atomic combinator $\mathbf{C}^{\gamma}$ in the definition of u^{β} become occurrences of the term $\lambda x^{\alpha}.\, \mathbf{C}^{\gamma}$, and this case can be treated as the previous one, since

$$\mathbf{K}^{\gamma\to(\alpha\to\gamma)} \mathbf{C}^{\gamma} \;\longrightarrow_{\beta}\; \lambda x^{\alpha}.\, \mathbf{C}^{\gamma}.$$

4. The occurrences of an application $u_1^{\gamma\to\delta}u_2^{\gamma}$ in the definition of u^{β} become occurrences of the term
$$\lambda x^{\alpha}.\, u_1^{\gamma\to\delta}u_2^{\gamma}.$$
The idea is to postulate an atomic combinator **S** that would reduce to it when applied to $\lambda x^{\alpha}.\, u_1^{\gamma\to\delta}$ and $\lambda x^{\alpha}.\, u_2^{\gamma}$, which we have by the induction hypothesis. We thus need a name for the binary function
$$(\lambda x^{\alpha}.\, u_1^{\gamma\to\delta}, \lambda x^{\alpha}.\, u_2^{\gamma}) \longmapsto \lambda x^{\alpha}.\, u_1^{\gamma\to\delta}u_2^{\gamma}.$$
Since x^{α} is bound both in the arguments and in the value, to avoid confusion we change its name to z^{α} in the latter by using the α-rule:
$$(\lambda x^{\alpha}.\, u_1^{\gamma\to\delta}, \lambda x^{\alpha}.\, u_2^{\gamma}) \longmapsto \lambda z^{\alpha}.\, (u_1^{\gamma\to\delta}u_2^{\gamma})[x^{\alpha} := z^{\alpha}].$$
We now know that **S** will be a binary function whose values are unary functions. Equivalently, using the convention on functions of many variables, we can think of it as a ternary function, as follows:
$$(\lambda x^{\alpha}.\, u_1^{\gamma\to\delta}, \lambda x^{\alpha}.\, u_2^{\gamma}, z^{\alpha}) \longmapsto (u_1^{\gamma\to\delta}u_2^{\gamma})[x^{\alpha} := z^{\alpha}].$$
By definition of substitution,
$$(u_1^{\gamma\to\delta}u_2^{\gamma})[x^{\alpha} := z^{\alpha}] = (u_1^{\gamma\to\delta}[x^{\alpha} := z^{\alpha}])(u_2^{\gamma}[x^{\alpha} := z^{\alpha}]).$$
Since, for $i = 1, 2$,
$$(\lambda x^{\alpha}.\, u_i)z^{\alpha} \longrightarrow_{\beta} u_i[x^{\alpha} := z^{\alpha}],$$
it is enough to define **S** as a ternary function that first applies its first two arguments separately to the third, and then applies the two results. This is precisely how **S** is defined in 3 above.

We can now check that this intuition works:
$$\begin{array}{ll} & \mathbf{S}^{[\alpha\to(\gamma\to\delta)]\to[(\alpha\to\gamma)\to(\alpha\to\delta)]}(\lambda x^{\alpha}.\, u_1)(\lambda x^{\alpha}.\, u_2) \\ \longrightarrow_{\beta} & \lambda z^{\alpha}.\, [(\lambda x^{\alpha}.\, u_1)z^{\alpha}][(\lambda x^{\alpha}.\, u_2)z^{\alpha}] \\ \longrightarrow_{\beta} & \lambda z^{\alpha}.\, (u_1[x^{\alpha} := z^{\alpha}])(u_2[x^{\alpha} := z^{\alpha}]) \\ \longrightarrow_{\beta} & \lambda z^{\alpha}.\, (u_1u_2)[x^{\alpha} := z^{\alpha}] \\ \longrightarrow_{\beta} & \lambda x^{\alpha}.\, u_1u_2, \end{array}$$
where the next to last step holds by definition of substitution (in the case of an application), and the last step holds by the α-rule. □

The names of the three families of atomic combinators introduced above reflect their definitions: **I** stands for *identity*; **K** indicates (to German speakers) the fact that its values are *constant* functions; and **S** recalls that its definition involves a *substitution* (plus composition).

It should be noted that the family of atomic combinators $\mathbf{I}^{\alpha\to\alpha}$, one for each α, is actually derivable from the other two, as $\mathbf{SKK}$. More precisely, as:

$$(\mathbf{S}^{\{\alpha\to[(\alpha\to\alpha)\to\alpha]\}\to\{[\alpha\to(\alpha\to\alpha)]\to(\alpha\to\alpha)\}}\mathbf{K}^{\alpha\to[(\alpha\to\alpha)\to\alpha]})\mathbf{K}^{\alpha\to(\alpha\to\alpha)}.$$

Thus only two families of atomic combinators are enough to synthesize every typed λ-term by composition alone, starting from the variables.

In a nutshell, the proof just given amounts to the following inductive translation of λ-terms into combinatorial terms, where we drop types for readability:

$$\begin{array}{rcl}\lambda x.\, x & = & \mathbf{I} \\ \lambda x.\, y & = & \mathbf{K}y \\ \lambda x.\, \mathbf{C} & = & \mathbf{KC} \quad (\mathbf{C} = \mathbf{K} \text{ or } \mathbf{S}) \\ \lambda x.\, u_1u_2 & = & \mathbf{S}(\lambda x.\, u_1)(\lambda x.\, u_2).\end{array}$$

It should be noted that: globally, $\mathbf{K}$ and $\mathbf{S}$ are not the only possible choice of atomic combinators (and historically, not even the first ones); locally, the previous translation is not the only possible one, in terms of $\mathbf{K}$ and $\mathbf{S}$. Different choices of atomic combinators and/or translations are made with an eye to issues of efficiency and complexity, discussed in Peyton Jones [1987].

Exercise 8.2.2 *A single family of atomic combinators is enough to synthesize every typed λ-term.* (Meredith and Prior [1963], Barendregt) (Hint: we can look for a combinator $\mathbf{C}$ such that $\mathbf{CS} = \mathbf{K}$. Since $\mathbf{S}$ is a function of two variables, it is enough to let

$$\mathbf{C} = \lambda x.\, xabc,$$

with a, b and c combinators to be determined. Then

$$\mathbf{CS} = \mathbf{S}abc = (ac)(bc),$$

and to get $\mathbf{K}$ as a result it is enough to let $a = c = \mathbf{K}$. It remains to determine b in such a way that some iteration of $\mathbf{C}$ produces $\mathbf{S}$. If

$$\mathbf{C} = \lambda x.\, x\mathbf{K}b\mathbf{K},$$

then automatically

$$\mathbf{CC} = \mathbf{KK} \qquad \text{and} \qquad \mathbf{C}(\mathbf{CC}) = b.$$

It is thus enough to let

$$\mathbf{C} = \lambda x.\, x\mathbf{KSK}$$

to have $\mathbf{C}(\mathbf{CC}) = \mathbf{S}$ and $\mathbf{CS} = \mathbf{K}$.

Obviously, we have to consider different typed versions of $\mathbf{K}$ in the definition of $\mathbf{C}$, and different typed versions of $\mathbf{C}$ in the various iterations. Thus, if

$$\mathbf{C}_i = \lambda x.\, x\mathbf{K}_i\mathbf{S}_i\mathbf{K}_i^*,$$

then $\mathbf{C}_1\mathbf{C}_2 = \mathbf{K}_2^*\mathbf{K}_1^*$, $\mathbf{C}_3(\mathbf{C}_1\mathbf{C}_2) = \mathbf{S}_3$, and $\mathbf{C}_4\mathbf{S}_3 = \mathbf{K}_4$.

Symmetrically, we can look for a combinator $\mathbf{D}$ such that $\mathbf{DK} = \mathbf{S}$. Since $\mathbf{K}$ is a function of three variables, it is enough to let

$$\mathbf{D} = \lambda x.\, xab,$$

with a and b combinators to be determined. Then

$$\mathbf{DK} = \mathbf{K}ab = a,$$

and to get as result $\mathbf{S}$ it is enough to let $a = \mathbf{S}$. It remains to determine b in such a way that some iteration of $\mathbf{D}$ produces $\mathbf{K}$. If

$$\mathbf{D} = \lambda x.\, x\mathbf{S}b,$$

then automatically

$$\mathbf{DD} = \lambda z.\, (bz)(bbz).$$

To get $\mathbf{K}$ directly it is enough to define b as a function of two variables that applied to z and anything else produces $\lambda y.\, z$, so that

$$\lambda z.\, (bz)(bbz) = \lambda z.\, \lambda y.\, z = \mathbf{K}.$$

Then it is enough to let

$$b = \lambda zxy.\, z \qquad \text{and } \mathbf{D} = \lambda x.\, x\mathbf{S}b$$

to have $\mathbf{DD} = \mathbf{K}$ and $\mathbf{DK} = \mathbf{S}$.)

Having shown how combinatorial terms are actually sufficient to define all λ-terms, we can take a further step and develop a **Typed Theory of Combinators** independently of the Typed Lambda Calculus, and as an alternative approach to it. Briefly, this goes as follows.

The notion of a *typed combinator* is defined inductively, as in 8.1.1, using in a first approximation only the two constants $\mathbf{K}$ and $\mathbf{S}$ (in a second approximation, the presence of other combinator constants distinguishes the *pure* from an *applied* theory of combinators):

1. The constants $\mathbf{K}^{\gamma\to(\alpha\to\gamma)}$ and $\mathbf{S}^{[\alpha\to(\gamma\to\delta)]\to[(\alpha\to\gamma)\to(\alpha\to\delta)]}$ are combinators of the indicated types.

2. If $\mathbf{C}^{\alpha\to\beta}$ and $\mathbf{D}^{\alpha}$ are combinators of type $\alpha \to \beta$ and α, respectively, then $(\mathbf{C}^{\alpha\to\beta}\mathbf{D}^{\alpha})^{\beta}$ is a combinator of type β.

Notice how λ-abstraction has disappeared, and with it any notion of bound variable (together with the relative α-rule).

To increase readability some parentheses or types can be omitted, when no confusion arises. We will use the letters x, y, z, … for variables, and $\mathbf{C}$, $\mathbf{D}$, $\mathbf{E}$, … for combinators.

The notion of *combinatorial β-reducibility* $\longrightarrow_{c\beta}$ is defined in analogy with β-reducibility, by replacing the β-rule with its two instances needed to give the constants $\mathbf{K}$ and $\mathbf{S}$ their appropriate operational meaning. Precisely, $\longrightarrow_{c\beta}$ is the reflexive and transitive closure of the single step reducibility $\longrightarrow_{1c\beta}$, defined inductively as:

$$\mathbf{K}^{\gamma\to(\alpha\to\gamma)}\mathbf{C}^{\gamma}\mathbf{D}^{\alpha} \longrightarrow_{1c\beta} \mathbf{C}^{\gamma}$$

$$\mathbf{S}^{[\alpha\to(\gamma\to\delta)]\to[(\alpha\to\gamma)\to(\alpha\to\delta)]}\mathbf{C}^{\alpha\to(\gamma\to\delta)}\mathbf{D}^{\alpha\to\gamma}\mathbf{E}^{\alpha} \longrightarrow_{1c\beta} (\mathbf{CE})^{\gamma\to\delta}(\mathbf{DE})^{\gamma}$$

$$\frac{\mathbf{C}_1^{\alpha\to\beta} \longrightarrow_{1c\beta} \mathbf{C}_2^{\alpha\to\beta}}{\mathbf{C}_1^{\alpha\to\beta}\mathbf{D}^{\alpha} \longrightarrow_{1c\beta} \mathbf{C}_2^{\alpha\to\beta}\mathbf{D}^{\alpha}}$$

$$\frac{\mathbf{D}_1^{\alpha} \longrightarrow_{1c\beta} \mathbf{D}_2^{\alpha}}{\mathbf{C}^{\alpha\to\beta}\mathbf{D}_1^{\alpha} \longrightarrow_{1c\beta} \mathbf{C}^{\alpha\to\beta}\mathbf{D}_2^{\alpha}.}$$

The notion of *combinatorial β-equality* $=_{c\beta}$ is defined as the symmetric and transitive closure of $\longrightarrow_{c\beta}$.

A combinator is in *β-normal form* if no β-reduction is possible inside it or, equivalently, if it does not contain any subcombinator of the form **KCD** or **SCDE**.

The proof of 8.2.1 shows how to define, for every λ-term t, a combinator t_c that β-reduces to it. On the other hand, any combinator $\mathbf{C}$ can be directly translated into a λ-term $\mathbf{C}_{\lambda}$, by plugging in the definitions of $\mathbf{S}$ and $\mathbf{K}$ given in 8.2.1. The question naturally arises of how the two translations are related, in particular whether one is the inverse of the other. This will be aswered in the affirmative by 8.6.9 and 8.6.10, in the presence of extensionality rules.

8.3 Existence of Normal Forms

The terms of the Typed Lambda Calculus can be considered as descriptions of objects. Among them the ones with nonatomic types, i.e. with types of the form $\alpha \to \beta$, can be considered as descriptions of functions. The β-rule describes an atomic step in the computation of a value of a given function for a given argument. The β-reduction procedure allows the performance of this atomic step any finite number of times, inside given terms. If a β-reduction is a computation, then its values are reached when the computation terminates, i.e. when a term in normal form is obtained.

For practical purposes, terms in normal form can thus be taken not only as *describing* the objects of the Typed Lambda Calculus, as all other terms, but as *being* the objects themselves. If this sounds odd, we can go back again to the example of the calculus of polynomial expressions. Here we do have the natural numbers as intended objects in mind, but their definition, although a philosophically interesting problem, is neither needed nor useful in the calculus. For practical purposes, a number can just be thought of as coinciding with a polynomial expression in normal form, i.e. its decimal representation, since the calculus is going to stop there anyway. And this practical attitude is all that is needed to carry out computations.

Weak Normalization

We now show that every typed λ-term has a meaning, i.e. it denotes an object in the sense just described. Otherwise said, *every λ-term has a normal form.* A corollary of this is that the Typed Lambda Calculus is *a calculus of total functions*, i.e. the value of a function for a given argument always exists.

It would not be very useful to have a nonconstructive proof of the existence of normal forms, without an algorithm to obtain them. The proof of the next theorem is constructive, and it can be abstractly seen as consisting of two parts: the description of a *computational strategy*, specifying how reductions should performed, and a *termination proof* showing that the strategy works, i.e. that it always produces a normal form.

Theorem 8.3.1 Weak Normalization (Turing [1942], Curry and Feys [1958]) *Every typed λ-term can be reduced to a normal term, by means of an appropriate sequence of β-reductions.*

Proof. Notice that the reduction of a redex $(\lambda x^\alpha .\, u^\beta)^{\alpha\to\beta} v^\alpha$ inside a given term t into the redex $u^\beta[x^\alpha := v^\alpha]$ can have the following two bad effects:

- it can increase the total number of redexes, since it substitutes v^α (and hence all redexes occurring in it) for every free occurrence of x^α in u^β, and there may be many such occurrences;

- it can introduce new redexes, in two different ways:

  - if $\alpha = \gamma \to \delta$ and $v^{\gamma\to\delta}$ is of the form $\lambda z^\gamma .\, v_1^\delta$, by turning certain subterms of t of the form $x^{\gamma\to\delta} w^\gamma$ (precisely, the ones in the scope of the redex in question) into redexes $(\lambda z^\gamma .\, v_1^\delta) w^\gamma$
  - if $\beta = \gamma \to \delta$ and $u^{\gamma\to\delta}[x^\alpha := v^\alpha]$ is of the form $\lambda z^\gamma .\, u_1^\delta$, by turning certain subterms of t of the form $((\lambda x^\alpha .\, u^{\gamma\to\delta}) v^\alpha) w_1^\gamma$ into redexes $(\lambda z^\gamma .\, u_1^\delta) w_1^\gamma$.

The main observation is that the second obstacle is not traumatic, since the new redexes possibly introduced are of complexity lower than the one being eliminated. The appropriate measure of complexity for $(\lambda x^\alpha .\, u^\beta)^{\alpha\to\beta} v^\alpha$ is in this case the **degree** of the type $\alpha \to \beta$, defined inductively as follows:

- atomic types, i.e. type letters, have degree 0

- the degree of $\alpha \to \beta$ is 1 plus the greatest of the degree of α and β.

The idea of the normalizing procedure is thus to eliminate, at every step, a redex of greatest degree, until all of them have been disposed of. The first obstacle is overcome by choosing, at every step, *a redex* $(\lambda x^\alpha . \, u^\beta)v^\alpha$ *of greatest degree, such that in* v^α *no redex of greatest degree occurs* (so that only the number of redexes of degree smaller than the greatest one can be increased).[5]

By so doing, at every step we eliminate one redex of greatest degree, and do not introduce new ones of the same degree. Once the last redex of greatest degree has been eliminated, we attack the ones of the next greatest degree (whose number, in the meantime, may have greatly increased), and so on, until all redexes have been eliminated. □

Formally, the proof of the Weak Normalization Theorem is by so-called ω^2-*induction*, i.e. induction on pairs of natural numbers (a, b) lexicographically ordered by

$$(a, b) \prec (a', b') \;\Leftrightarrow\; (a < a') \vee (a = a' \wedge b < b').$$

At every step the pair

$$(\text{greatest degree}, \text{number of redexes with greatest degree})$$

strictly decreases in the ordering $\prec$. In other words, either the greatest degree decreases, or it remains the same but the number of redexes with greatest degree decreases by one.

Syntactical proof of Strong Normalization

The Weak Normalization Theorem showed that a clever strategy of reductions always produces a normal form. The Strong Normalization Theorem shows that we do not have to be clever in choosing our reduction strategy: any order would eventually do.

Theorem 8.3.2 Strong Normalization (Hinatani [1966], Hinata [1967], Sanchis [1967], Shoenfield [1967], Tait [1967], Dragalin [1968]). *For every typed* λ*-term* t^α*, there is no infinite sequence of reductions starting from* t^α*.*

Corollary 8.3.3 *For every typed* λ*-term* t^α*, the tree of all possible reductions starting from* t^α *is finite.*

[5] Such a redex can be found by the following inductive procedure. First choose a redex of greatest degree. If it contains no redex of greatest degree, then stop. Otherwise, choose one of the redexes of greates degree contained in it, and start again. After finitely many steps we must hit a redex of greatest degree containing no redexes of greatest degree.

Proof of Corollary. The tree of all possible reductions starting from a given term is finitely branching, since only finitely many reductions are possible inside a given term. Then the corollary is a trivial consequence of 8.3.2 and of *König's Lemma*, according to which a finitely branching tree with no infinite branch is finite.

We prove the contrapositive of König's Lemma: if a finitely branching tree is infinite, then it has an infinite branch. Let n_0 be the root of the tree. Consider the nodes of level 1, i.e. those immediately following the root n_0. By definition of finitely branching tree, there are only finitely many. And since the tree is infinite, at least one of these nodes has infinitely many successors. Choose one node n_1 of level 1 with infinitely many successors, and consider the nodes of level 2 which immediately follow n_1. As above, there must be one node n_2 following n_1 and having infinitely many successors. By continuing the procedure, we get a sequence of nodes $n_0, n_1, n_2, \ldots$, each a successor of the previous ones. Moreover, each of these nodes has infinitely many successors, and thus the procedure never stops. In particular, the n_i's define an infinite branch. $\Box$

We turn now to a proof of the Strong Normalization Theorem. We call a term *strongly normalizable* if there is no infinite sequence of reductions starting from it. The theorem then says that every term is strongly normalizable.

Since we are dealing with typed λ-terms, to prove the theorem we can proceed by induction either on types or on terms.

If we do *induction on types*, we can easily deal with the case of nonatomic types. Given $t^{\alpha\to\beta}$, to be able to apply the induction hypothesis we need an object of a lower type, and a natural choice is $t^{\alpha\to\beta}x^{\alpha}$, with x^{α} any variable of type α. Any infinite sequence of reductions in $t^{\alpha\to\beta}$ would produce an infinite sequence of reductions in $t^{\alpha\to\beta}x^{\alpha}$, which is impossible by the induction hypothesis.

The real problem is how to deal with terms of atomic type. We will attempt a proof by induction on terms in a moment, but first we stop to notice that we have actually proved that *if $\mathcal{C}$ is any class of terms*

1. *containing only strongly normalizable terms of atomic type*

2. *closed under application to variables,*

then $\mathcal{C}$ contains only strongly normalizable terms.

We can now try to fill the gap in the first proof, by attempting to prove that every term of atomic type is strongly normalizable. In fact, we may just as well try to prove the full theorem by *induction on terms*. If t is a variable, then no reduction is possible inside it, and thus t is strongly normalizable. If t is a λ-abstraction $\lambda x.\, u$, then the only possible reductions are inside u, which is strongly normalizable by the induction hypothesis.

The real problem is how to prove that the application uv of two strongly normalizable terms u and v is still strongly normalizable. The difficulty is due to the fact that u might be a λ-abstraction, and thus in uv there might be reductions not retraceable to reductions inside u and v. Then the induction hypothesis cannot be applied.

We can now merge the two approaches, in two steps. First, we strengthen the condition on $\mathcal{C}$ above and request closure not only under applications to variables, but under every application. Second, we turn such a condition into a definition, by changing an implication ('if $t^{\alpha\to\beta} \in \mathcal{C}$, then $t^{\alpha\to\beta}u^{\alpha} \in \mathcal{C}$, for any $u^{\alpha} \in \mathcal{C}$') into an equivalence.

Then the first proof, by induction on types, will show that all terms in $\mathcal{C}$ are strongly normalizable, and the second proof, by induction on terms, will show that every term is in $\mathcal{C}$. The two proofs together will then show that every term is strongly normalizable.

Definition 8.3.4 (Tait [1967]) *The class $\mathcal{C}$ of* **computable terms** *is defined by induction on types, as follows:*

1. *for α atomic,*
$$t^{\alpha} \in \mathcal{C} \;\Leftrightarrow\; t^{\alpha} \text{ is strongly normalizable}$$

2. *for $\alpha \to \beta$,*
$$t^{\alpha\to\beta} \in \mathcal{C} \;\Leftrightarrow\; (\forall u^{\alpha} \in \mathcal{C})(t^{\alpha\to\beta}u^{\alpha} \in \mathcal{C}).$$

The first proof above would be immediately reproducible, if we knew that $\mathcal{C}$ contains all variables. Actually, since variables are only needed to lower the type, it would be enough to know $\mathcal{C}$ *is not trivial*, i.e. that it contains at least one term (precisely, a variable) for any given type.

To try to prove that $x^{\alpha} \in \mathcal{C}$, we can proceed by induction on types. If α is atomic, then $x^{\alpha} \in \mathcal{C}$ because a variable is strongly normalizable. But $x^{\alpha\to\beta} \in \mathcal{C}$ if and only if $x^{\alpha\to\beta}u^{\alpha} \in \mathcal{C}$, for any $u^{\alpha} \in \mathcal{C}$. Now $x^{\alpha\to\beta}u^{\alpha}$ does have a type β smaller than $\alpha \to \beta$, but the induction does not apply: it tells us that *variables* of type smaller than $\alpha \to \beta$ are in $\mathcal{C}$, but $x^{\alpha\to\beta}u^{\alpha}$ is not a variable.

We thus need to strengthen the induction hypothesis, and attempt to prove not that x^{α} is in $\mathcal{C}$, but rather that $x^{\alpha}u_1^{\alpha_1}\cdots u_n^{\alpha_n}$ is, for any $u_1,\ldots,u_n \in \mathcal{C}$ with the appropriate types. We proceed, as before, by induction on types. Now the inductive step has been fixed, but we find a difficulty at atomic types: if $x^{\alpha\to\beta}u_1^{\alpha_1}\cdots u_n^{\alpha_n}$ has an atomic type, then it is in $\mathcal{C}$ if and only if it is strongly normalizable. For this we need to know that the u_i are strongly normalizable, while we only know that they are in $\mathcal{C}$. Notice how the fact that 'terms in $\mathcal{C}$ are strongly normalizable' is exactly what we are trying to prove!

The atomic step would go through if we only wanted to prove that $x^{\alpha\to\beta}u_1^{\alpha_1}\cdots u_n^{\alpha_n}$ is in $\mathcal{C}$ when the u_i are strongly normalizable. Is this weaker statement enough to proceed inductively? It would be, if at types $\alpha \to \beta$ we knew that terms u in $\mathcal{C}$ with smaller types are strongly normalizable. This would be given by the induction hypothesis, if we were trying to prove *simultaneously* that 'terms in $\mathcal{C}$ are strongly normalizable'. This finally produces the needed conditions.

Proposition 8.3.5 (Tait [1967]) *By simultaneous induction on types we can prove:*

1. *if $u_1, \ldots, u_n$ are strongly normalizable, then $(xu_1 \cdots u_n)^\alpha \in \mathcal{C}$*
2. *if $t^\alpha \in \mathcal{C}$, then t^α is strongly normalizable.*

Proof. We first consider atomic types. If $xu_1 \cdots u_n$ has an atomic type, then it is in $\mathcal{C}$ if and only if it is strongly normalizable. But since x is a variable and

$$xu_1 \cdots u_n = (\cdots(xu_1)\cdots u_n),$$

the only possible reductions are inside the u_i, which are strongly normalizable. Thus $xu_1 \cdots u_n$ is strongly normalizable.

If t^α has atomic type, then it is strongly normalizable if it is in $\mathcal{C}$, by definition of $\mathcal{C}$.

We now consider arrow types. If $xu_1 \cdots u_n$ has type $\alpha \to \beta$, then it is in $\mathcal{C}$ if and only if $xu_1 \cdots u_n u^\alpha$ is, for any $u^\alpha \in \mathcal{C}$. By the induction hypothesis 2, u^α is strongly normalizable. Since the induction hypothesis 1 works for any number of terms, in particular for $n+1$, then $xu_1 \cdots u_n u$ is in $\mathcal{C}$.

Given $t^{\alpha\to\beta} \in \mathcal{C}$, by the induction hypothesis 1 (with $n = 0$) $x^\alpha \in \mathcal{C}$. Then $t^{\alpha\to\beta}x^\alpha$ is in $\mathcal{C}$ by definition, and is strongly normalizable by the induction hypothesis 2 (having type β). Then $t^{\alpha\to\beta}$ is strongly normalizable too, since any infinite sequence of reductions inside it would produce an infinite sequence of reductions inside $t^{\alpha\to\beta}x^\alpha$. □

We turn now to the second half of the proof, and try to prove that every term is in $\mathcal{C}$, by induction on terms. We already know from 8.3.5.1 (with $n = 0$) that all variables are in $\mathcal{C}$. If u and v are in $\mathcal{C}$ then so is uv, since $\mathcal{C}$ is closed under application by definition. We are left with the case of a term $(\lambda x^\alpha.\, u^\beta)^{\alpha\to\beta}$, which is in $\mathcal{C}$ if and only if $(\lambda x^\alpha.\, u^\beta)^{\alpha\to\beta}v^\alpha$ is, for any $v^\alpha \in \mathcal{C}$. Now $(\lambda x^\alpha.\, u^\beta)^{\alpha\to\beta}v^\alpha$ only *reduces* to $u^\beta[x^\alpha := v^\alpha]$, and we find two difficulties. First, even if we knew that the latter reduct is in $\mathcal{C}$, we would need to prove that the original redex is. Second, the induction hypothesis is too weak to show that $u^\beta[x^\alpha := v^\alpha]$ is in $\mathcal{C}$, since it only ensures that u^β is in $\mathcal{C}$.

The second difficulty can easily be remedied, by strengthening the induction hypothesis. The next lemma takes care of the first difficulty.

Proposition 8.3.6 (Tait [1967]) *For any term u^β and v^α and any variable x^α,*

$$u^\beta[x^\alpha := v^\alpha] \in \mathcal{C} \wedge v^\alpha \in \mathcal{C} \Rightarrow (\lambda x^\alpha .\, u^\beta)v^\alpha \in \mathcal{C}.$$

Proof. By the inductive definition of types, the general form of the type β of $(\lambda x^\alpha .\, u^\beta)v^\alpha$ is

$$\alpha_1 \to (\alpha_2 \to \cdots (\alpha_n \to \alpha_{n+1}) \cdots),$$

with α_{n+1} atomic. By repeatedly using the definition of $\mathcal{C}$, we have to prove that

$$u_1^{\alpha_1}, \ldots, u_n^{\alpha_n} \in \mathcal{C} \Rightarrow (\lambda x^\alpha .\, u^\beta)v^\alpha u_1^{\alpha_1} \cdots u_n^{\alpha_n} \in \mathcal{C}.$$

The right-hand-side has atomic type α_{n+1}, and it is thus in $\mathcal{C}$ if and only if it is strongly normalizable.

Any reduction in $(\lambda x^\alpha .\, u^\beta)v^\alpha u_1^{\alpha_1} \cdots u_n^{\alpha_n}$ either eliminates the first λ and produces $u^\beta[x^\alpha := v^\alpha]u_1^{\alpha_1} \cdots u_n^{\alpha_n}$, or it can also be performed inside the latter term (if it is made inside u^β or some $u_i^{\alpha_i}$), or it is a reduction in v^α.[6] In all cases, only finitely many such reductions can be performed, since both $u^\beta[x^\alpha := v^\alpha]u_1^{\alpha_1} \cdots u_n^{\alpha_n}$ and v^α are in $\mathcal{C}$ (the first one by closure of $\mathcal{C}$ under composition and the hypothesis $u^\beta[x^\alpha := v^\alpha] \in \mathcal{C}$, the second one by the hypothesis $v^\alpha \in \mathcal{C}$), and by 8.3.5 every term in $\mathcal{C}$ is strongly normalizable. □

The next result provides the final step, with the appropriate strengthening of the induction hypothesis.

Proposition 8.3.7 (Tait [1967]) *By induction on terms, we can prove that for any t^α, $x_i^{\alpha_i}$ and $u_i^{\alpha_i}$:*

$$u_1^{\alpha_1}, \ldots, u_n^{\alpha_n} \in \mathcal{C} \Rightarrow t^\alpha[x_1^{\alpha_1} := u_1^{\alpha_1}, \ldots, x_n^{\alpha_n} := u_n^{\alpha_n}] \in \mathcal{C},$$

where the substitutions are performed simultaneously.

Proof. To improve readability, we write

$$[\vec{x} := \vec{u}] \qquad \text{for} \qquad [x_1^{\alpha_1} := u_1^{\alpha_1}, \ldots, x_n^{\alpha_n} := u_n^{\alpha_n}].$$

If t is a variable, then either it is $x_i^{\alpha_i}$, in which case the result of the substitution is $u_i^{\alpha_i}$, which is in $\mathcal{C}$ by the induction hypothesis; or it is a variable different from all $x_i^{\alpha_i}$'s, in which case the substitution has no effect and t is in $\mathcal{C}$, being a variable.

If t is an application $t_1 t_2$, then

$$(t_1 t_2)[\vec{x} := \vec{u}] = (t_1[\vec{x} := \vec{u}])(t_2[\vec{x} := \vec{u}])$$

[6] Note that the this case is superfluous if x^α occurs free in u^β, but is needed otherwise because then $u^\beta[x^\alpha := v^\alpha]$ does not contain occurrences of v^α, and reductions inside v^α cannot necessarily be performed inside $u^\beta[x^\alpha := v^\alpha]u_1^{\alpha_1} \cdots u_n^{\alpha_n}$.

by definition of substitution. By the induction hypothesis, both $t_1[\vec{x} := \vec{u}]$ and $t_2[\vec{x} := \vec{u}]$ are in $\mathcal{C}$, and hence so is the above right-hand-side, by closure of $\mathcal{C}$ under composition. Then so is the left-hand-side.

If t is a λ-abstraction $\lambda y^\alpha.\, t_1$, where y may be supposed to be different from all the $\vec{x}$ by the α-rule, then

$$(\lambda y.\, t_1)[\vec{x} := \vec{u}] = \lambda y.\, (t_1[\vec{x} := \vec{u}])$$

by definition of substitution. The right-hand-side is in $\mathcal{C}$ if and only if, for any $v^\alpha \in \mathcal{C}$,

$$[\lambda y^\alpha.\, (t_1[\vec{x} := \vec{u}])]v^\alpha \in \mathcal{C}.$$

This reduces to

$$t_1[\vec{x} := \vec{u}, x^\alpha := v^\alpha],$$

in which the substitutions can be supposed to be performed simultaneously (by possibly renaming the bound variable y^α). Since the induction hypothesis applies to t_1, this term is in $\mathcal{C}$. Hence so is $[\lambda y^\alpha.\, (t_1[\vec{x} := \vec{u}])]v^\alpha$, by 8.3.6. □

Corollary 8.3.8 Strong Normalization. *Every typed λ-term t^α is strongly normalizable.*

Proof. By 8.3.7, $t^\alpha \in \mathcal{C}$. By 8.3.5, t^α is strongly normalizable. □

Regarding the complexity of the proof just given, we should notice that the original definition 8.3.4 of the class $\mathcal{C}$ is quite complicated: it introduces a new quantifier at every type level, and it thus roughly requires n quantifiers for terms whose type degree is at most n. In particular, the definition of the full class $\mathcal{C}$ seems to require 'infinitely many quantifiers'. Technically stated, and modulo a coding of λ-terms as numbers, the definition of $\mathcal{C}$ given in the proof is not arithmetical, and *the proof given above is not formalizable in First-Order Arithmetic*. Of course, the proof shows that $\mathcal{C}$ is actually the class of *all* terms, which is certainly easily definable, but this is *after* the facts.

On the proofs of Strong Normalization ⋆

The proof of Strong Normalization given above consists of two separate parts: one proved by induction on types (8.3.5), and the other by induction on terms (8.3.7). Its interest lies in the fact that it generalizes to a number of other type systems, and it is thus very useful. However, it does not provide an intuitive picture of what is going on.

More perspicuous is the proof given by Howard [1970] in the style of 8.3.1, which consists in assigning ordinals to terms in such a way that they strictly decrease at every reduction step of any reduction procedure (not only of a particular one, as in 8.3.1). We provide a semantical version of such a proof in 10.5, in which the ordinal assignment is replaced by an interpretation of terms on elements of appropriate well-founded sets. Incidentally, such a proof will be formalizable in First-Order Arithmetic.

For those only interested in the result itself, without concerns of generality or perspicuity, we now give a short and direct *alternative proof* by induction on terms, as follows. The cases of a variable and of a λ-abstraction are trivial. Consider then an application uv of two strongly normalizable terms u and v, and let $h(u)$ and $h(v)$ be the height of the reduction trees starting from u and v, respectively. It is enough to prove, by induction on $h(u) + h(v)$, that every term which is one reduction step away from uv is strongly normalizable.

If the reduction step is made inside u or v, then one of $h(u)$ and $h(v)$ decreases, and the induction hypothesis applies. We are thus reduced to the case in which u is a λ-abstraction, and the reduction step is of the form

$$(\lambda x.\, u_1)v \longrightarrow_{1\beta} u_1[x := v].$$

This is taken care of by the following proposition.

Proposition 8.3.9 (Nederpelt [1973]) *If t^β and v^α are strongly normalizable, then so is $t^\beta[x^\alpha := v^\alpha]$.*

Proof. By a triple induction on:

1. the height $h(t) + h(v)$
2. the type α
3. the term t,

we prove that every term which is one reduction step away from $t[x := v]$ is strongly normalizable. There are two cases:

- If the reduction step is made inside t or v and produces a term of the form $t_1[x := v]$ or $t[x := v_1]$, then the induction hypothesis 1 applies because one of $h(t)$ or $h(v)$ decreases.
- Otherwise, the reduced redex is of the form $(\lambda y.\, s)w$, and it must have been created by the substitution of $\lambda y.\, s$ for x. By replacing $(\lambda y.\, s)w$ by a fresh variable z, we obtain a term t' such that

  $$t[x := v] = t'[z := (\lambda y.\, s)w].$$

  Since $\lambda y.\, s$ has the same type as x, z must have a smaller type. By the induction hypothesis 2, to prove that every term which is one reduction step away from $t[x := v]$ is strongly normalizable, it is enough to show that both t' and $(\lambda y.\, s)w$ are strongly normalizable. There are two cases:

- If $t[x = v]$ is different from $(\lambda y.\, s)w$, then both t' and $(\lambda y.\, s)w$ are obtained by substituting v for x in subterms of t, and they are strongly normalizable by the induction hypothesis 3.

- If $t[x = v]$ is equal to $(\lambda y.\, s)w$, then the only possible reductions from it must be inside s, w or $s[y := w]$. The terms s and w are strongly normalizable by the hypothesis, because they are subterms of v and t. The term $s[y := w]$ is strongly normalizable by the induction hypothesis 2, because y has type smaller than x. □

8.4 Uniqueness of Normal Forms

Having disposed of the problem of existence of normal forms, we now turn to the complementary problem of uniqueness.

The idea behind the proof of uniqueness is quite simple: if we start from a given term t and β-reduce it into different ways, thus obtaining two terms t_1 and t_2, then permuting the steps (i.e. doing in t_1 what we did in t to get t_2, and in t_2 what we did in t to get t_1) should produce the same term t^*. This is the so-called *Diamond Property*, pictured as follows:

$$\begin{tikzcd} t \arrow[r] \arrow[d] & t_1 \arrow[d] \\ t_2 \arrow[r] & t^* \end{tikzcd}$$

From this, uniqueness of normal forms follows immediately: if t_1 and t_2 are already in normal form, they cannot be reduced and thus

$$t_1 \;=\; t^* \;=\; t_2.$$

The obvious strategy to prove the Diamond Property is to proceed by induction, by finding an appropriate way of splitting $\rightarrow_\beta$ in blocks of steps $\Rightarrow$, and proving the Diamond Property for $\Rightarrow$, i.e.

$$\begin{tikzcd} u \arrow[r, Rightarrow] \arrow[d, Rightarrow] & u_1 \arrow[d, Rightarrow] \\ u_2 \arrow[r, Rightarrow] & u^*. \end{tikzcd}$$

Then we can proceed inductively on the number of applications of $\Rightarrow$ from $t = t_{00}$

to $t_1 = t_{0n}$ and $t_2 = t_{m0}$, in a way that we will refer to as ‘diagram chasing’:

$$\begin{tikzcd}
t_{00} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_{01} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_{02} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_{03} \arrow[d, phantom, "\vdots"] & \cdots & t_{0n} \arrow[d, Rightarrow] \\
t_{10} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_{11} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_{12} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & \cdots \arrow[d, phantom, "\vdots"] & \cdots & t_{1n} \arrow[d, Rightarrow] \\
t_{20} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_{21} \arrow[r, Rightarrow] \arrow[d, Rightarrow] & \cdots \arrow[r, phantom, "\cdots"] \arrow[d, phantom, "\vdots"] & \cdots \arrow[d, phantom, "\vdots"] & \cdots & t_{2n} \arrow[d, Rightarrow] \\
t_{30} \arrow[r, Rightarrow] \arrow[d, phantom, "\vdots"] & \cdots \arrow[r, phantom, "\cdots"] \arrow[d, phantom, "\vdots"] & \cdots \arrow[r, phantom, "\cdots"] \arrow[d, phantom, "\vdots"] & \cdots \arrow[d, phantom, "\vdots"] & \cdots & t_{3n} \arrow[d, phantom, "\vdots"] \\
t_{m0} \arrow[r, Rightarrow] & t_{m1} \arrow[r, Rightarrow] & t_{m2} \arrow[r, Rightarrow] & t_{m3} & \cdots & t_{mn}.
\end{tikzcd}$$

The Weak Diamond Property

The first splitting of $\rightarrow_{\beta}$ that comes to mind is of course $\rightarrow_{1\beta}$. Unfortunately, the Diamond Property fails for β-reductions consisting of at most one application of the β-rule. What does hold is the so-called *Weak Diamond Property*: not quite as symmetric as we would like, but still both a first step toward the full result, and sufficient to prove uniqueness of normal forms (when coupled with Strong Normalization, as we will see in 8.4.5).[7]

Proposition 8.4.1 Weak Diamond Property (Newmann [1942]) *If t_1 and t_2 are terms obtained from t by at most one application of the β-rule, then they reduce to a common term t^* by a finite number (depending on t) of applications of the β-rule.*

Proof. We proceed by induction on t. Since types play no role in the proof, we do not bother to write them down.

If t is a variable, then no application of the β-rule is possible inside t. Hence $t_1 = t_2 = t^*$.

If $t = \lambda x.\, u$, then applications of the β-rule are possible only inside u. Hence $t_1 = \lambda x.\, u_1$ and $t_2 = \lambda x.\, u_2$, where u_1 and u_2 are obtained from u by at most one application of the β-rule. By the induction hypothesis, they reduce to a common term u^*, and thus both t_1 and t_2 reduce to $t^* = \lambda x.\, u^*$.

If $t = uv$, then there are a number of possibilities:

[7] Notice that the lack of symmetry in the Weak Diamond Property for $\longrightarrow_{1\beta}$ forbids its use in a proof of the full Diamond Property for $\longrightarrow_{\beta}$ by induction on the number of β-reduction. One β-reduction is in general replaced by many, and the induction produces an expanding diagram that does not converge.

- t_1 and t_2 are obtained by at most one application of the β-rule inside u. Then $t_1 = u_1 v$ and $t_2 = u_2 v$. By the induction hypothesis, u_1 and u_2 reduce to a common term u^*, and thus both t_1 and t_2 reduce to $t^* = u^* v$.

- t_1 and t_2 are obtained by at most one application of the β-rule inside v. This case is similar to the previous one.

- t_1 and t_2 are obtained by at most one application of the β-rule inside u and v, separately. Thus e.g. $t_1 = u_1 v$ and $t_2 = uv_2$, and they reduce to $t^* = u_1 v_2$.

- One of t_1 and t_2 is obtained by at most one application of the β-rule inside u, and the other is obtained by one application across u and v. Thus e.g. $u = (\lambda x.\, u')$, $t = (\lambda x.\, u')v$, $t_1 = u'[x := v]$ and $t_2 = (\lambda x.\, u'_2)v$. Then t_1 and t_2 reduce to the common term $t^* = u'_2[x := v]$. This is trivial for t_2, and requires a simple proof for t_1 (see 8.4.2).

- One of t_1 and t_2 is obtained by at most one application of the β-rule inside v, and the other is obtained by one application across u and v. Thus e.g. $u = (\lambda x.\, u')$, $t = (\lambda x.\, u')v$, $t_1 = u'[x := v]$ and $t_2 = (\lambda x.\, u')v_2$. Then t_1 and t_2 reduce to the common term $t^* = u'[x := v_2]$. This is trivial for t_2, and requires a simple proof for t_1 (see 8.4.4). □

We have left two loose strings in the proof just finished, which we now tie.

Proposition 8.4.2 *For any u, u^* and t,*

$$\frac{u \longrightarrow_{1\beta} u^*}{u[x := t] \longrightarrow_{1\beta} u^*[x := t].}$$

Proof. By induction on $u \longrightarrow_{1\beta} u^*$.

1. $(\lambda y.\, u_1)v \longrightarrow_{1\beta} u_1[y := v]$
   By definition of substitution:

$$\begin{array}{rcl} u[x := t] & = & (\lambda y.\, u_1[x := t])(v[x := t]) \\ u^*[x := t] & = & (u_1[y := v])[x := t]. \end{array}$$

   Now
$$(\lambda y.\, u_1[x := t])(v[x := t]) \longrightarrow_{1\beta} (u_1[x := t])[y := v[x := t]].$$

   By the following property of substitution:

$$(u_1[x := t])[y := v[x := t]] = (u_1[y := v])[x := t]$$

   the result follows.

The property just quoted amounts to saying that if we first substitute v for y in u_1, and then t for x in the result, then we obtain the same result as if we first substitute t for x in both u_1 and v separately, and then substitute $v[x := t]$ for y in u_1. This is quite intuitive, and can easily be proved using the definition of substitution (see 8.4.3).

2. $u_1 v \longrightarrow_{1\beta} u_2 v$, with $u_1 \longrightarrow_{1\beta} u_2$
By definition of substitution:

$$\begin{array}{rcccl} u[x := t] & = & (u_1 v)[x := t] & = & (u_1[x := t])(v[x := t]) \\ u^*[x := t] & = & (u_2 v)[x := t] & = & (u_2[x := t])(v[x := t]). \end{array}$$

And

$$\dfrac{\dfrac{u_1 \longrightarrow_{1\beta} u_2}{u_1[x := t] \longrightarrow_{1\beta} u_2[x := t]}}{(u_1[x := t])(v[x := t]) \longrightarrow_{1\beta} (u_2[x := t])(v[x := t])}$$

by hypothesis, induction hypothesis, and 8.2 of 8.1.4.

3. $u_1 v_1 \longrightarrow_{1\beta} u_1 v_2$, with $v_1 \longrightarrow_{1\beta} v_2$
Similar to part 2, using 8.3 of 8.1.4.

4. $\lambda y.\, u_1 \longrightarrow_{1\beta} \lambda y.\, u_2$, with $u_1 \longrightarrow_{1\beta} u_2$
By definition of substitution:

$$\begin{array}{rcccl} u[x := t] & = & (\lambda y.\, u_1)[x := t] & = & \lambda y.\, (u_1[x := t]) \\ u^*[x := t] & = & (\lambda y.\, u_2)[x := t] & = & \lambda y.\, (u_2[x := t]). \end{array}$$

And

$$\dfrac{\dfrac{u_1 \longrightarrow_{1\beta} u_2}{u_1[x := t] \longrightarrow_{1\beta} u_2[x := t]}}{\lambda y.\, (u_1[x := t]) \longrightarrow_{1\beta} \lambda y.\, (u_2[x := t])}$$

by hypothesis, induction hypothesis, and 8.4 of 8.1.4. □

Exercise 8.4.3 Substitution Property. *If y is not free in t, then*

$$(u[x := t])[y := v[x := t]] = (u[y := v])[x := t].$$

(Hint: by induction on u. The crucial case is when $u = x$. Then

$$(u[x := t])[y := v[x := t]] = t[y := v[x := t]] \qquad \text{and} \qquad (u[y := v])[x := t] = t,$$

and the two sides are equal if y is not free in t. Obviously, the latter condition is crucial here, but it does no harm in the proof of 8.4.2, where y is bound by λ, and thus the α-rule allows us to suppose that y is not free in t.)

Proposition 8.4.4 *For any t, t^* and u,*

$$\frac{t \longrightarrow_{1\beta} t^*}{u[x := t] \longrightarrow_\beta u[x := t^*].}$$

Proof. By induction on u.

1. If $u = x$, then $u[x := t] = t$ and $u[x = t^*] = t^*$, and $t \longrightarrow_\beta t^*$ is given by hypothesis.

   If u is a variable different from x, then $u[x := t] = u[x := t^*] = u$, and $u \longrightarrow_\beta u$ by definition of $\longrightarrow_\beta$.

2. If $u = u_1 u_2$, then

   $$\begin{array}{lclcl} u[x := t] & = & (u_1 u_2)[x := t] & = & (u_1[x := t])(u_2[x := t]) \\ u[x := t^*] & = & (u_1 u_2)[x := t^*] & = & (u_1[x := t^*])(u_2[x := t^*]) \end{array}$$

   by definition of substitution. Notice that

   $$\begin{array}{c} t \longrightarrow_{1\beta} t^* \\ \hline u_1[x := t] \longrightarrow_\beta u_1[x := t^*] \\ \hline (u_1[x := t])(u_2[x := t]) \longrightarrow_\beta (u_1[x := t^*])(u_2[x := t]) \end{array}$$

   by hypothesis, induction hypothesis, and 8.2 of 8.1.4. Similarly,

   $$\begin{array}{c} t \longrightarrow_{1\beta} t^* \\ \hline u_2[x := t] \longrightarrow_\beta u_2[x := t^*] \\ \hline (u_1[x := t^*])(u_2[x := t]) \longrightarrow_\beta (u_1[x := t^*])(u_2[x := t^*]) \end{array}$$

   by hypothesis, induction hypothesis, and 8.3 of 8.1.4. By transitivity, i.e. 8.7 of 8.1.5, we then have

   $$(u_1[x := t])(u_2[x := t]) \longrightarrow_\beta (u_1[x := t^*])(u_2[x := t^*]).$$

3. If $u = \lambda y.\, u_1$, then

   $$\begin{array}{lclcl} u[x := t] & = & (\lambda y.\, u_1)[x := t] & = & \lambda y.\, (u_1[x := t]) \\ u[x := t^*] & = & (\lambda y.\, u_1)[x := t^*] & = & \lambda y.\, (u_1[x := t^*]) \end{array}$$

   by definition of substitution. And

   $$\begin{array}{c} t \longrightarrow_{1\beta} t^* \\ \hline u_1[x := t] \longrightarrow_\beta u_1[x := t^*] \\ \hline \lambda y.\, (u_1[x := t]) \longrightarrow_\beta \lambda y.\, (u_1[x := t^*]) \end{array}$$

   by hypothesis, induction hypothesis, and 8.4 of 8.1.4. □

The last result and a look at the proofs of 8.4.1, 8.4.2 and 8.4.4 show that we 'almost' proved the Diamond Property for $\longrightarrow_{1\beta}$. What we missed is the strong form of 8.4.4, namely

$$\frac{t \longrightarrow_{1\beta} t^*}{u[x := t] \longrightarrow_{1\beta} u[x := t^*].}$$

What went wrong in the proof of 8.4.4 is case 2, where we needed to apply transitivity of $\longrightarrow_{\beta}$, which obviously does not hold for $\longrightarrow_{1\beta}$.

Actually, it is not only the proposed proof that fails, but result itself. To step from $u[x := t]$ to $u[x := t^*]$ we need to perform the step from t to t^* in every occurrence of t inside $u[x := t]$, and the number of such occurrences may be arbitrarily high. Indeed, every free occurrence of x in u produces an occurrence of t in $u[x := t]$, and there may be any number of such free occurrences.

In particular, no limitation on the *number of applications* of the β-rule is going to produce the symmetric inversion of steps needed for the Diamond Property. What we would need is the possibility of making *parallel independent reductions*, e.g. of making the same reduction inside v in all occurrences of v inside u. This suggests the definition of a notion of parallel reduction, which will be introduced in the next section.

For our present purposes, however, the Weak Diamond Property is already enough for a proof of the following result, as noticed by Huet [1980].

Theorem 8.4.5 Uniqueness of Normal Forms. *Every typed λ-term has exactly one normal form.*

Proof. The existence part follows from the Weak Normalization Theorem 8.3.1. The uniqueness part is proved by induction on the height of the tree of all possible reductions, and this requires the Strong Normalization Theorem 8.3.2 (to know that the tree has finite height).

Given t, let t_1 and t_2 be two terms in normal form to which it reduces, and consider the first step t'_1 and t'_2 in each of the two reductions. Since they are obtained from t by one application of the β-rule, by the Weak Diamond Property 8.4.1 they reduce to a common term t^*, that we may suppose to be in normal form (otherwise, by the Normalization Theorem we just reduce it to a normal form).

Consider now the tree of all possible reductions from t'_1. Since t reduces to t'_1, this tree is a subtree of the tree of all possible reductions from t, and it has a height smaller than it by at least one, since one reduction is required to go from t to t'_1. Then, by the induction hypothesis, any two normal forms of t'_1 are equal. In particular, $t_1 = t^*$.

Similarly, by working on the tree relative to t'_2, we get $t_2 = t^*$. But then $t_1 = t_2$, since they are both equal to t^*. □

Equality again

From existence and uniqueness of normal forms we can deduce an alternative characterization of β-equality.

Proposition 8.4.6 *Two typed terms are β-equal if and only if they have the same normal form.*

Proof. One direction has been proved in 8.1.8. For the opposite direction, suppose $u_1 =_\beta u_2$. By induction on 8.1.6, there are three cases:

- $u_1 \longrightarrow_\beta u_2$
  Then the normal form of u_2, which exists and is unique by 8.3.1 and 8.4.5, is also the normal form of u_1.

- $u_2 =_\beta u_1$
  By the induction hypothesis, u_2 and u_1 have the same normal form.

- $u_1 =_\beta t$ and $t =_\beta u_2$
  By the induction hypothesis, u_1 and t have the same normal form, and so do t and u_2. Then u_1 and u_2 have the same normal form. □

Corollary 8.4.7 *$=_\beta$ is decidable.*

Proof. By 8.4.6, to check whether u_1 and u_2 are equal it is enough to compute their normal forms, and see if they are the same. □

The complexity of the previous decision procedure depends on the complexity of its two steps, i.e. normalization and comparison of normal forms. The latter is trivial. The complexity of the former will be determined in 11.4.1 and 11.4.2.

8.5 The Church-Rosser Theorem ⋆

The work in the previous section was motivated by the search for a proof of the following result.

Theorem 8.5.1 Diamond Property (Church and Rosser [1936]) *If t_1 and t_2 are terms obtained from t by $\longrightarrow_\beta$, then there is a term t^* which can be obtained from t_1 and t_2 by $\longrightarrow_\beta$.*

If we really only wanted a proof, then we could already give two:

- *by Weak Normalization and Uniqueness of Normal Forms*
  If t_1 and t_2 are obtained from t by $\longrightarrow_\beta$, it is enough to let t^* be the normal form of t_1 and t_2, which is uniquely determined because it also the normal form of t.

- *by Strong Normalization and Weak Diamond Property*
  We proceed as in 8.4.5. Let t_1' and t_2' be the first steps of the reductions of t to t_1 and t_2. By the Weak Diamond Property, they reduce to a common term t'. By the induction hypothesis applied to t_1', t_1 and t' reduce to a common term u_1. By the induction hypothesis applied to t_2', t' and t_2 reduce to a common term u_2. By the induction hypothesis applied to t', u_1 and u_2 reduce to a common term t^*.

$$\begin{tikzcd}
 & & t_1 \arrow[dr, "\beta"] & & \\
 & t_1' \arrow[ur, "\beta"'] \arrow[dr, "\beta"'] & & u_1 \arrow[dr, "\beta"'] & \\
t \arrow[ur, "1\beta"'] \arrow[dr, "1\beta"'] & & t' \arrow[ur, "\beta"'] \arrow[dr, "\beta"] & & t^* \\
 & t_2' \arrow[ur, "\beta"'] \arrow[dr, "\beta"'] & & u_2 \arrow[ur, "\beta"'] & \\
 & & t_2 \arrow[ur, "\beta"'] & &
\end{tikzcd}$$

However, both of these proofs use a form of normalization, and they would thus not work for the Untyped Lambda Calculus.

The goal of this section is to give a simplified version of the original proof due to Church and Rosser, which avoids any use of normalization.

Parallel reductions

The proof of the Weak Diamond Property showed quite clearly what is needed for a notion of reduction to have the Diamond Property: it should commute with the term formation rules for the trivial cases to go through, and it should avoid the trouble expressed by the failure of the rule

$$\frac{v \longrightarrow_{1\beta} v^*}{u[x := v] \longrightarrow_{1\beta} u[x := v^*],}$$

i.e. allow for parallel reductions. Such a notion can be defined inductively, following the pattern of 8.1.4.

Since notations will be pretty heavy and types play no role in the proof, as in 8.4.1 we do not bother to write them down.

Definition 8.5.2 Parallel Reducibility (Tait, Martin-Löf) *The reducibility $\Longrightarrow$ is defined inductively by the following clauses:*

$$u \Longrightarrow u \tag{8.11}$$

$$\frac{u \Longrightarrow u_1 \quad v \Longrightarrow v_1}{uv \Longrightarrow u_1 v_1} \tag{8.12}$$

$$\frac{u \Longrightarrow u_1}{\lambda x.\, u \Longrightarrow \lambda x.\, u_1} \tag{8.13}$$

$$\frac{u \Longrightarrow u_1 \quad v \Longrightarrow v_1}{(\lambda x.\, u)v \Longrightarrow u_1[x := v_1],} \tag{8.14}$$

where the first clause can be thought of as an axiom, and the remaining ones as deduction rules.

As a special case of equation 8.14 we have

$$\frac{u \Longrightarrow u \quad v \Longrightarrow v}{(\lambda x.\, u)v \Longrightarrow u[x := v],}$$

where the top line consists of axioms. Thus $\longrightarrow_{1\beta}$ implies $\Longrightarrow$. Inductively, we easily see that $\Longrightarrow$ implies $\longrightarrow_{\beta}$. Then $\longrightarrow_{\beta}$ *is the transitive closure of* $\Longrightarrow$, since it is the transitive closure of $\longrightarrow_{1\beta}$.

But $\Longrightarrow$ *is not transitive* itself. For example,

$$(\lambda x^{\alpha\to\alpha}.\, x^{\alpha\to\alpha} y^{\alpha})(\lambda z^{\alpha}.\, z^{\alpha}) \Longrightarrow (\lambda z^{\alpha}.\, z^{\alpha}) y^{\alpha} \Longrightarrow y^{\alpha}$$

holds because each step is a single application of β-reduction, but

$$(\lambda x^{\alpha\to\alpha}.\, x^{\alpha\to\alpha} y^{\alpha})(\lambda z^{\alpha}.\, z^{\alpha}) \Longrightarrow y^{\alpha},$$

fails because of the following crucial properties of $\Longrightarrow$:

1. *if* $\lambda x.\, u \Longrightarrow v$, *then* $v = \lambda x.\, u_1$, *with* $u \Longrightarrow u_1$
2. *if* $uv \Longrightarrow t$, *then one of the following holds:*
   - $t = u_1 v_1$ *with* $u \Longrightarrow u_1$ *and* $v \Longrightarrow v_1$
   - $u = \lambda x.\, u'$ *and* $t = u'_1[x := v_1]$, *with* $u' \Longrightarrow u'_1$ *and* $v \Longrightarrow v_1$.

In other words, the lack of transitivity allows us to retrace our steps back. In logical terminology, *the deduction system associated with* $\Longrightarrow$ *is cut-free.*

Before proving the Diamond Property for $\Longrightarrow$ we prove the substitution property that held in some special cases (8.4.2), but failed in general for $\longrightarrow_{1\beta}$. The result is a simultaneous version of 8.4.2 and 8.4.4, and the proof combines the proofs of 8.4.2 (by induction on $\longrightarrow_{1\beta}$) and 8.4.4 (by induction on u).

Proposition 8.5.3 (Tait, Martin-Löf) *The following is a derived rule for* $\Longrightarrow$*:*

$$\frac{u \Longrightarrow u^* \qquad t \Longrightarrow t^*}{u[x := t] \Longrightarrow u^*[x := t^*].}$$

Proof. By induction on $u \Longrightarrow u^*$.

1. $u \Longrightarrow u$
   We want

   $$\frac{t \Longrightarrow t^*}{u[x := t] \Longrightarrow u[x := t^*],}$$

   which we prove by induction on u.

   (a) If $u = x$, then $u[x := t] = t$ and $u[x = t^*] = t^*$, and $t \Longrightarrow t^*$ is given by hypothesis.

   If u is a variable different from x, then $u[x := t] = u[x := t^*] = u$, and $u \Longrightarrow u$ is an axiom.

   (b) If $u = u_1 u_2$, then

   $$\begin{array}{lclcl} u[x := t] & = & (u_1 u_2)[x := t] & = & (u_1[x := t])(u_2[x := t]) \\ u[x := t^*] & = & (u_1 u_2)[x := t^*] & = & (u_1[x := t^*])(u_2[x := t^*]) \end{array}$$

   by definition of substitution. And

   $$\frac{\dfrac{t \Longrightarrow t^*}{u_1[x := t] \Longrightarrow u_1[x := t^*]} \qquad \dfrac{t \Longrightarrow t^*}{u_2[x := t] \Longrightarrow u_2[x := t^*]}}{(u_1[x := t])(u_2[x := t]) \Longrightarrow (u_1[x := t^*])(u_2[x := t^*])}$$

   by hypothesis, induction hypothesis, and 8.12 of 8.5.2.

   (c) If $u = \lambda y.\, u_1$, then

   $$\begin{array}{lclcl} u[x := t] & = & (\lambda y.\, u_1)[x := t] & = & \lambda y.\, (u_1[x := t]) \\ u[x := t^*] & = & (\lambda y.\, u_1)[x := t^*] & = & \lambda y.\, (u_1[x := t^*]) \end{array}$$

   by definition of substitution. And

   $$\frac{\dfrac{t \Longrightarrow t^*}{u_1[x := t] \Longrightarrow u_1[x := t^*]}}{\lambda y.\, (u_1[x := t]) \Longrightarrow \lambda y.\, (u_1[x := t^*])}$$

   by hypothesis, induction hypothesis, and 8.13 of 8.5.2.

2. $u_1 v_1 \Longrightarrow u_2 v_2$, with $u_1 \Longrightarrow u_2$ and $v_1 \Longrightarrow v_2$
   By definition of substitution:

$$\begin{array}{rcccl} u[x := t] & = & (u_1 v_1)[x := t] & = & (u_1[x := t])(v_1[x := t]) \\ u^*[x := t^*] & = & (u_2 v_2)[x := t^*] & = & (u_2[x := t^*])(v_2[x := t^*]). \end{array}$$

   And

$$\dfrac{\dfrac{u_1 \Longrightarrow u_2 \quad t \Longrightarrow t^*}{u_1[x := t] \Longrightarrow u_2[x := t^*]} \qquad \dfrac{v_1 \Longrightarrow v_2 \quad t \Longrightarrow t^*}{v_1[x := t] \Longrightarrow v_2[x := t^*]}}{(u_1[x := t])(v_1[x := t]) \Longrightarrow (u_2[x := t^*])(v_2[x := t^*])}$$

   by hypothesis, induction hypothesis, and 8.12 of 8.5.2.

3. $\lambda y.\, u_1 \Longrightarrow \lambda y.\, u_2$, with $u_1 \Longrightarrow u_2$
   By definition of substitution:

$$\begin{array}{rcccl} u[x := t] & = & (\lambda y.\, u_1)[x := t] & = & \lambda y.\, (u_1[x := t]) \\ u^*[x := t^*] & = & (\lambda y.\, u_2)[x := t^*] & = & \lambda y.\, (u_2[x := t^*]). \end{array}$$

   And

$$\dfrac{\dfrac{u_1 \Longrightarrow u_2 \quad t \Longrightarrow t^*}{u_1[x := t] \Longrightarrow u_2[x := t^*]}}{\lambda y.\, (u_1[x := t]) \Longrightarrow \lambda y.\, (u_2[x := t^*])}$$

   by hypothesis, induction hypothesis, and 8.13 of 8.5.2.

4. $(\lambda y.\, u_1) v_1 \Longrightarrow u_2[y := v_2]$, with $u_1 \Longrightarrow u_2$ and $v_1 \Longrightarrow v_2$
   By definition of substitution:

$$\begin{array}{rcl} u[x := t] & = & (\lambda y.\, u_1[x := t])(v_1[x := t]) \\ u^*[x := t^*] & = & (u_2[y := v_2])[x := t^*]. \end{array}$$

   Notice that

$$\dfrac{\dfrac{u_1 \Longrightarrow u_2 \quad t \Longrightarrow t^*}{u_1[x := t] \Longrightarrow u_2[x := t^*]} \qquad \dfrac{v_1 \Longrightarrow v_2 \quad t \Longrightarrow t^*}{v_1[x := t] \Longrightarrow v_2[x := t^*]}}{(\lambda y.\, u_1[x := t])(v_1[x := t]) \Longrightarrow (u_2[x := t^*])[y := v_2[x := t^*]]}$$

   by hypothesis, induction hypothesis, and 8.14 of 8.5.2. Since

$$(u_2[x := t^*])[y := v_2[x := t^*]] = (u_2[y := v_2])[x := t^*]$$

   by the Substitution Property (8.4.3), the result follows. □

We can now prove the result that motivated the introduction of $\Longrightarrow$.

Proposition 8.5.4 Strong Diamond Property for $\Longrightarrow$ (Takahashi [1995]) *For any term t there is a term t^* such that if $t \Longrightarrow t_1$, then $t_1 \Longrightarrow t^*$.*

Proof. Intuitively, we define t^* by performing all possible reductions in t in parallel. Formally, we define t^* by induction on t, as follows:

$$t^* = \begin{cases} x & \text{if } t = x \\ u^* v^* & \text{if } t = uv \text{ and } u \text{ is not a } \lambda\text{-abstraction} \\ \lambda x.\, u^* & \text{if } t = \lambda x.\, u \\ u^*[x := v^*] & \text{if } t = (\lambda x.\, u)v. \end{cases}$$

We prove by induction on t that if $t \Longrightarrow t_1$, then $t_1 \Longrightarrow t^*$.

1. $t = x$
   If $t \Longrightarrow t_1$, it must be $t_1 = x$. Then
   $$t_1 = x \Longrightarrow x = t^*$$
   by 8.11 of 8.5.2.

2. $t = uv$, with u not a λ-abstraction
   If $t \Longrightarrow t_1$, it must be $t_1 = u_1 v_1$, with $u \Longrightarrow u_1$ and $v \Longrightarrow v_1$. Then
   $$\frac{u_1 \Longrightarrow u^* \qquad v_1 \Longrightarrow v^*}{t_1 = u_1 v_1 \Longrightarrow u^* v^* = t^*}$$
   by the induction hypothesis and 8.12 of 8.5.2.

3. $t = \lambda x.\, u$
   If $t \Longrightarrow t_1$, it must be $t_1 = \lambda x.\, u_1$, with $u \Longrightarrow u_1$. Then
   $$\frac{u_1 \Longrightarrow u^*}{t_1 = \lambda x.\, u_1 \Longrightarrow \lambda x.\, u^* = t^*}$$
   by the induction hypothesis and 8.13 of 8.5.2.

4. $t = (\lambda x.\, u)v$
   If $t \Longrightarrow t_1$, there are two subcases:

   (a) $t_1 = (\lambda x.\, u_1)v_1$, with $u \Longrightarrow u_1$ and $v \Longrightarrow v_1$
   Then
   $$\frac{u_1 \Longrightarrow u^* \qquad v_1 \Longrightarrow v^*}{t_1 = (\lambda x.\, u_1)v_1 \Longrightarrow u^*[x := v^*] = t^*}$$
   by the induction hypothesis and 8.14 of 8.5.2.

(b) $t_1 = u_1[x := v_1]$, with $u \Longrightarrow u_1$ and $v \Longrightarrow v_1$
Then

$$\frac{u_1 \Longrightarrow u^* \qquad v_1 \Longrightarrow v^*}{t_1 = u_1[x := v_1] \Longrightarrow u^*[x := v^*] = t^*}$$

by the induction hypothesis and 8.5.3. □

Corollary 8.5.5 Diamond Property for $\Longrightarrow$ (Tait, Martin-Löf) *If t_1 and t_2 are terms obtained from t by $\Longrightarrow$, then there is a term t^* which can be obtained from t_1 and t_2 by $\Longrightarrow$. Graphically,*

$$\begin{tikzcd} t \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_1 \arrow[d, Rightarrow] \\ t_2 \arrow[r, Rightarrow] & t^*. \end{tikzcd}$$

Theorem 8.5.6 Diamond Property for $\longrightarrow_\beta$ (Church and Rosser [1936]) *If t_1 and t_2 are terms obtained from t by $\longrightarrow_\beta$, then there is a term t^* which can be obtained from t_1 and t_2 by $\longrightarrow_\beta$. Graphically,*

$$\begin{tikzcd} t \arrow[r] \arrow[d] & t_1 \arrow[d] \\ t_2 \arrow[r] & t^*. \end{tikzcd}$$

Proof. By 'diagram chasing' (as at the beginning of this section), using the fact that $\longrightarrow_\beta$ can be split into a sequence of $\longrightarrow_{1\beta}$, and hence of $\Longrightarrow$. □

Corollary 8.5.7 Uniqueness of Normal Forms. *Every typed λ-term has exactly one normal form.*

Proof. If t_1 and t_2 are normal forms of t, by the previous corollary they reduce to a common term t^*. But since they are in normal form, no reduction is possible inside them, and thus $t_1 = t^* = t_2$. □

Normalization and the Church-Rosser Theorem ⋆

Note that the rules 8.5.2 defining $\Longrightarrow$ can be taken as the inductive definition of a class of reductions φ, with the following properties (where $u \Longrightarrow v$ is written as $v = \varphi(u)$):

$$\varphi(u) = u \tag{8.15}$$

$$\varphi(uv) = \varphi(u)\varphi(v) \tag{8.16}$$

$$\varphi(\lambda x.\, u) = \lambda x.\, \varphi(u) \tag{8.17}$$

$$\varphi((\lambda x.\, u)v) = \varphi(u)[x := \varphi(v)]. \tag{8.18}$$

Thus $u \Longrightarrow v$ means that there is a sequence of reductions that works inside out, i.e. from inner redexes to outer ones, and such a sequence is called a *minimal development*. We mark a certain number of λ's, and eliminate them by successively reducing minimal ones, i.e. ones with no inner λ's. If 8.15 above is replaced by

$$\varphi(x) = x,$$

then, by descending to the atomic terms, i.e. the variables, we get sequences called *complete minimal developments*.

The content of the proof of the Church-Rosser Theorem given above is thus similar to a *Weak Normalization Theorem*, showing that a particular reduction strategy produces the desired result.

An alternative proof of the Church-Rosser Theorem would be the following, similar to a *Strong Normalization Theorem*:

1. Show that for any term there are only finitely many possible developments. This is done by assigning indices to terms in such a way that they decrease along any development.

2. Show that all complete developments of a term w.r.t. a given set of λ's have the same value. This follows by considering any pair of selected redexes, and seeing what happens after they have been reduced in different orders (by cases on the various possibilities, i.e. on whether the redexes are disjoint or overlapping).

3. Show that the notion

   $$u \rightsquigarrow v \;\Leftrightarrow\; \text{there is a complete development from } u \text{ to } v$$

   has the Diamond Property. This follows from the two parts above: if $u \rightsquigarrow u_1$ and $u \rightsquigarrow u_2$, consider the sets R_1 and R_2 of redexes selected in u to go to u_1 and u_2, and complete the two developments (which were only complete w.r.t. R_1 and R_2, respectively) w.r.t. $R_1 \cup R_2$. The procedure terminates by part 1, and it produces the same result by part 2.

4. The Church-Rosser Theorem then follows by diagram chasing, since $\longrightarrow_\beta$ is the transitive closure of $\rightsquigarrow$.

Equality once more

From the Diamond Property we can deduce another characterization of β-equality which also works for the Untyped Lambda Calculus, unlike the previous characterization 8.4.6.

Proposition 8.5.8 *Two terms are β-equal if and only if they reduce to a common term.*

Proof. One direction has been proved in 8.1.7. For the opposite direction, suppose $u_1 =_\beta u_2$. By induction on 8.1.6, there are three cases:

- $u_1 \longrightarrow_\beta u_2$
  Then there is nothing to prove.

- $u_2 =_\beta u_1$
  By the induction hypothesis, u_2 and u_1 reduce to a common term.

- $u_1 =_\beta t$ and $t =_\beta u_2$
  By the induction hypothesis, u_1 and t reduce to a common term v_1, and t and u_2 reduce to a common term v_2. By the Diamond Property, v_1 and v_2 reduce to a common term t^*:

$$\begin{tikzcd}
u_1 \arrow[dr, "\beta"'] & & \\
& v_1 \arrow[dr, "\beta"'] & \\
t \arrow[ur, "\beta"'] \arrow[dr, "\beta"'] & & t^* \\
& v_2 \arrow[ur, "\beta"'] & \\
u_2 \arrow[ur, "\beta"'] & &
\end{tikzcd}$$

  Then u_1 and u_2 reduce to a common term t^*. □

8.6 Extensionality

Since the terms of the Typed Lambda Calculus code programs that compute functions, the problem arises of extensionally identifying intensionally different programs that compute the same function. In this section we introduce a number of equivalent solutions to the problem, and study the extensional version of the Typed Lambda Calculus thus obtained.

When needed to avoid confusion, we will explicitly refer to notions, symbols and results of the previous sections by attaching a β to them. For example, we will talk of β-normal forms, parallel β-reductions $\Rightarrow_\beta$ and strong β-normalization.

Extensional reductions

There are at least three possible ways of identifying different terms that compute the same function:

- the infinitary *ext-rule*, which is a version of Leibniz's principle of indiscernibles, identifies two terms u and v if they cannot be distinguished on the basis of their behavior on arbitrary terms t:

$$\frac{(\forall t)(ut =_\beta vt)}{u =_{ext} v}$$

- the finitary ς-*rule*, which is a different version of the previous one, identifies two terms u and v if they cannot be distinguished on the basis of their behaviour on an arbitrary variable x not occurring free in any of them:

$$\frac{ux =_\beta vx}{u =_\varsigma v}$$

- the finitary η-*rule* identifies a term u with the function $\lambda x.\, ux$ whose values are obtained by applying u to a generic argument x not occurring free in u:

$$\lambda x.\, ux =_\eta u.$$

The next proposition shows that, luckily, we do not have to choose among the various rules.

Proposition 8.6.1 *The rules ext, ς and η are equivalent in the Typed Lambda Calculus.*

Proof. ς follows from η:

$$\frac{\dfrac{ux =_\beta vx}{\lambda x.\, ux =_\beta \lambda x.\, vx}}{u =_\eta \lambda x.\, ux =_\beta \lambda x.\, vx =_\eta v.}$$

ext follows from ς:

$$\frac{\dfrac{(\forall t)(ut =_\beta vt)}{ux =_\beta vx}}{u =_\varsigma v.}$$

η follows from *ext* because, since $(\lambda x.\, ux)t =_\beta ut$ holds for any t by the β-rule, then $\lambda x.\, ux =_{ext} u$. □

In particular, the infinitary *ext*-rule can be equationally presented by the η-rule, to which we confine in the following. The η-rule has been stated in the form of an equality rule, but it can easily be turned into a reduction rule as follows.

Definition 8.6.2 $\boldsymbol{\eta}$-rule. *Given a term u and a variable x not occurring free in it, we can step from $\lambda x.\, ux$ (called an **$\boldsymbol{\eta}$-redex**) to u (called an **$\boldsymbol{\eta}$-reduct**). We write*

$$\lambda x.\, ux \longrightarrow_{1\eta} u$$

to state that one step of the η-rule has been applied to the left hand side to produce the right hand side.

The following definition is the analogue of 8.1.4, 8.1.5 and 8.1.6.

Definition 8.6.3 $\boldsymbol{\eta}$-Reducibility and $\boldsymbol{\eta}$-Equality. *The reducibility $\longrightarrow_{1\eta}$ is defined inductively, by replacing in 8.1.4 the first clause by the following:*

$$\lambda x.\, ux \longrightarrow_{1\eta} u,$$

when x is not free in u.

The reducibility $\longrightarrow_\eta$ is the reflexive and transitive closure of $\longrightarrow_{1\eta}$.

The relation $=_\eta$ is the symmetric and transitive closure of $\longrightarrow_\eta$.

The reducibility relation $\longrightarrow_\eta$ is not particularly interesting on its own. For example, if we say that a term is in **$\boldsymbol{\eta}$-normal form** if no application of the η-rule is possible in it, then it is immediate to notice that *given a term, any sequence of applications of η-reductions to it will eventually produce a term in η-normal form* (**Strong $\boldsymbol{\eta}$-Normalization**), *that is independent of the chosen sequence of η-reductions* (**Uniqueness of $\boldsymbol{\eta}$-Normal Forms**). The first result is a trivial consequence of the observation that any application of the η-rule decreases the length of the term to which it is applied. The second result follows, as in 8.4.5, from the first and the Weak Diamond Property for η-reduction (the latter requires only the trivial part of the proof of 8.4.1).

It is thus more interesting to consider the η-rule not as a replacement of the β-rule, as above, but as a supplement to it.

Definition 8.6.4 $\boldsymbol{\beta\eta}$-Reducibility and $\boldsymbol{\beta\eta}$-Equality. *The reducibility $\longrightarrow_{1\beta\eta}$ is defined inductively, by adding to 8.1.4 the following clause:*

$$\lambda x.\, ux \longrightarrow_{1\eta} u,$$

when x is not free in u.

The reducibility $\longrightarrow_{\beta\eta}$ is the reflexive and transitive closure of $\longrightarrow_{1\beta\eta}$.

The relation $=_{\beta\eta}$ is the symmetric and transitive closure of $\longrightarrow_{\beta\eta}$.

Notice that $\beta\eta$-equality is coarser than pure β-equality, since e.g. $\lambda x.\, yx \neq_\beta y$, because the two sides are in β-normal form and different, but $\lambda x.\, yx =_{\beta\eta} y$ by definition.

As above, we say that a term is in ***βη*-normal form** if no application of the β-rule or η-rule is possible in it. In the following we prove that *given a term, any sequence of applications of $\beta\eta$-reductions to it will eventually produce a term in $\beta\eta$-normal form* (**Strong *βη*-Normalization**, 8.6.11), *that is independent of the chosen sequence of $\beta\eta$-reductions* (**Uniqueness of *βη*-Normal Forms**, 8.6.16).

Combinators ⋆

Of the three extensionality rules for the Typed Lambda Calculus introduced above, two admit obvious analogues for the Typed Theory of Combinators, namely:

- the infinitary *ext-rule*, which identifies two combinators **D** and **E** if they cannot be distinguished on the basis of their behavior on arbitrary combinators **C**:
$$\frac{(\forall \mathbf{C})(\mathbf{DC} =_{c\beta} \mathbf{EC})}{\mathbf{D} =_{ext} \mathbf{E}}$$
- the finitary ς-*rule*, which identifies two combinators **D** and **E** if they cannot be distinguished on the basis of their behaviour on an arbitrary variable x not occurring in any of them:
$$\frac{\mathbf{D}x =_{c\beta} \mathbf{E}x}{\mathbf{D} =_{\varsigma} \mathbf{E}.}$$

Proposition 8.6.5 *The rules ext and ς are equivalent in the Typed Theory of Combinators.*

Proof. To show that ς follows from *ext*, suppose $\mathbf{D}x =_{c\beta} \mathbf{E}x$ for some x not occurring in **D** or **E**. By substitution, $\mathbf{DC} =_{c\beta} \mathbf{EC}$ for any **C**. Then $\mathbf{D} =_{ext} \mathbf{E}$.

To show that *ext* follows from ς, suppose $(\forall \mathbf{C})(\mathbf{DC} =_{c\beta} \mathbf{EC})$. Then $\mathbf{D}x =_{c\beta} \mathbf{E}x$ for some x not occurring in **D** or **E**. Thus $\mathbf{D} =_{\varsigma} \mathbf{E}$. □

By adding any of the two equivalent extensionality rules to the theory of $c\beta$-equality, we obtain a theory of ***cβη*-equality**. Curry [1930] has proved that it is possible to equationally present such a theory by means of a *finite* set of equations. This is a better result than the one achieved above for the theory of $\beta\eta$-equality in the Typed Lambda Calculus, since the η-rule is only an *infinite* schema of equations (one for each term).

However, since our emphasis is on the extensional Typed Lambda Calculus, we do not develop here the extensional Typed Theory of Combinators, and content ourselves with showing that it possible to mutually translate one into the other. In one direction the translation is trivial, and it simply translates the atomic combinators in the natural way.

Definition 8.6.6 Lambda Translation of Combinators. *Given a combinator* $\mathbf{C}$*, we inductively translate it into a* λ*-term* $\mathbf{C}_\lambda$ *as follows:*

$$\mathbf{C}_\lambda = \begin{cases} x & \textit{if } \mathbf{C} = x \\ \lambda x.\, x & \textit{if } \mathbf{C} = \mathbf{I} \\ \lambda xy.\, x & \textit{if } \mathbf{C} = \mathbf{K} \\ \lambda xyz.\, (xz)(yz) & \textit{if } \mathbf{C} = \mathbf{S} \\ \mathbf{D}_\lambda \mathbf{E}_\lambda & \textit{if } \mathbf{C} = \mathbf{DE}. \end{cases}$$

In the other direction the translation is subtler, since it has to replace every possible occurrence of a λ-abstraction by a combinator.

Definition 8.6.7 Combinatorial Translation of Lambda Terms (Curry and Feys [1958]) *Given a* λ*-term* t*, we inductively translate it into a combinator* t_c *as follows:*

$$t_c = \begin{cases} x & \textit{if } t = x \\ u_c v_c & \textit{if } t = uv \\ \mathbf{I} & \textit{if } t = \lambda x.\, x \\ \mathbf{K}y & \textit{if } t = \lambda x.\, y \\ \mathbf{KC} & \textit{if } t = \lambda x.\, \mathbf{C} \qquad (\mathbf{C} = \mathbf{I},\, \mathbf{K} \textit{ or } \mathbf{S}) \\ \mathbf{S}(\lambda x.\, u)_c(\lambda x.\, v)_c & \textit{if } t = \lambda x.\, uv \\ (\lambda x.\, u_c)_c & \textit{if } t = \lambda x.\, u \textit{ and } u \textit{ is a } \lambda\textit{-abstraction.} \end{cases}$$

Before proving that the two translations are inverse one of the other, we state a pair of lemmas that will make it easier to deal with the crucial case of λ-abstraction.

Proposition 8.6.8 *For any combinator* $\mathbf{C}$*:*

1. $(\lambda x.\, \mathbf{C})_c x =_{c\beta} \mathbf{C}$
2. $(\lambda x.\, \mathbf{C})_{c\lambda} =_{\beta\eta} \lambda x.\, \mathbf{C}_\lambda$.

Proof. Part 1 is proved by induction on $\mathbf{C}$:

1. if $\mathbf{C} = x$, then
$$(\lambda x.\, x)_c x = \mathbf{I}x =_{c\beta} x.$$
2. if $\mathbf{C} = y$, then
$$(\lambda x.\, y)_c x = \mathbf{K}yx =_{c\beta} y.$$

3. if $\mathbf{C} = \mathbf{I}$, $\mathbf{K}$ or $\mathbf{S}$, then
$$(\lambda x.\, \mathbf{C})_c x = \mathbf{K}\mathbf{C}x =_{c\beta} \mathbf{C}.$$

4. if $\mathbf{C} = \mathbf{D}\mathbf{E}$, then
$$(\lambda x.\, \mathbf{D}\mathbf{E})_c x = \mathbf{S}(\lambda x.\, \mathbf{D})_c(\lambda x.\, \mathbf{E})_c x =_{c\beta} ((\lambda x.\, \mathbf{D})_c x)((\lambda x.\, \mathbf{E})_c x) =_{c\beta} \mathbf{D}\mathbf{E}$$
by the induction hypothesis.

Part 2 is also proved induction on $\mathbf{C}$:

1. if $\mathbf{C} = x$, then
$$(\lambda x.\, x)_{c\lambda} = \mathbf{I}_\lambda = \lambda x.\, x = \lambda x.\, x_\lambda.$$

2. if $\mathbf{C} = y$, then
$$(\lambda x.\, y)_{c\lambda} = (\mathbf{K}y)_\lambda = \mathbf{K}_\lambda y_\lambda = (\lambda zx.\, z)y_\lambda =_\beta \lambda x.\, y_\lambda.$$

3. if $\mathbf{C} = \mathbf{I}$, $\mathbf{K}$ or $\mathbf{S}$, then
$$(\lambda x.\, \mathbf{C})_{c\lambda} = (\mathbf{K}\mathbf{C})_\lambda = \mathbf{K}_\lambda \mathbf{C}_\lambda = (\lambda zx.\, z)\mathbf{C}_\lambda =_\beta \lambda x.\, \mathbf{C}_\lambda.$$

4. if $\mathbf{C} = \mathbf{D}\mathbf{E}$, then
$$\begin{array}{lll} (\lambda x.\, \mathbf{D}\mathbf{E})_{c\lambda} z & = & (\mathbf{S}(\lambda x.\, \mathbf{D})_c(\lambda x.\, \mathbf{E})_c)_\lambda z \\ & = & \mathbf{S}_\lambda(\lambda x.\, \mathbf{D})_{c\lambda}(\lambda x.\, \mathbf{E})_{c\lambda} z \\ & =_{\beta\eta} & \mathbf{S}_\lambda(\lambda x.\, \mathbf{D}_\lambda)(\lambda x.\, \mathbf{E}_\lambda) z \qquad \text{by the induction hypothesis} \\ & =_\beta & ((\lambda x.\, \mathbf{D}_\lambda)z)((\lambda x.\, \mathbf{E}_\lambda)z) \\ & =_\beta & (\mathbf{D}_\lambda[x := z])(\mathbf{E}_\lambda[x := z]) \\ & = & (\mathbf{D}_\lambda \mathbf{E}_\lambda)[x := z] \\ & = & (\mathbf{D}\mathbf{E})_\lambda[x := z] \\ & =_\beta & (\lambda x.\, (\mathbf{D}\mathbf{E})_\lambda)z. \end{array}$$

By extensionality, $(\lambda x.\, \mathbf{D}\mathbf{E})_{c\lambda} =_{\beta\eta} \lambda x.\, (\mathbf{D}\mathbf{E})_\lambda$. □

The next proposition shows that if we translate a combinator into a λ-term and the latter back into a combinator, we get an extensional identity.

Proposition 8.6.9 *For any combinator* $\mathbf{C}$*,* $(\mathbf{C}_\lambda)_c =_{c\beta\eta} \mathbf{C}$*.*

Proof. We proceed by induction on $\mathbf{C}$.

- If $\mathbf{C} = x$, then
$$(x_\lambda)_c = x_c = x.$$

- If $\mathbf{C} = \mathbf{I}$, then
$$(\mathbf{I}_\lambda)_c = (\lambda x.\, x)_c = \mathbf{I}.$$

- If $\mathbf{C} = \mathbf{K}$, then
$$(\mathbf{K}_\lambda)_c xy = (\lambda xy.x)_c xy = (\lambda x.\,(\lambda y.\, x)_c)_c xy =_{c\beta} (\lambda y.\, x)_c y =_{c\beta} x =_{c\beta} \mathbf{K}xy$$
by 8.6.8.1 twice. By extensionality, $(\mathbf{K}_\lambda)_c =_{c\beta\eta} \mathbf{K}$.

- If $\mathbf{C} = \mathbf{S}$, then
$$\begin{array}{rcl} (\mathbf{S}_\lambda)_c xyz & = & (\lambda xyz.\,(xz)(yz))_c xyz \\ & = & (\lambda x.\,(\lambda yz.\,(xz)(yz))_c)_c xyz \\ & =_{c\beta} & (\lambda yz.\,(xz)(yz))_c yz \\ & = & (\lambda y.\,(\lambda z.\,(xz)(yz))_c)_c yz \\ & =_{c\beta} & (\lambda z.\,(xz)(yz))_c z \\ & =_{c\beta} & (xz)(yz) \\ & =_{c\beta} & \mathbf{S}xyz \end{array}$$
by 8.6.8.1 three times. By extensionality, $(\mathbf{S}_\lambda)_c =_{c\beta\eta} \mathbf{S}$.

- If $\mathbf{C} = \mathbf{DE}$, then
$$((\mathbf{DE})_\lambda)_c = (\mathbf{D}_\lambda \mathbf{E}_\lambda)_c = (\mathbf{D}_\lambda)_c(\mathbf{E}_\lambda)_c =_{c\beta\eta} \mathbf{DE}$$
by the induction hypothesis. □

Similarly, the next proposition shows that if we translate a λ-term into a combinator and the latter back into a λ-term, we get an extensional identity.

Proposition 8.6.10 *For any λ-term t, $(t_c)_\lambda =_{\beta\eta} t$.*

Proof. We proceed by induction on t.

- If $t = x$, then
$$(x_c)_\lambda = x_\lambda = x.$$

- If $t = uv$, then
$$(uv)_{c\lambda} = (u_c v_c)_\lambda = u_{c\lambda} v_{c\lambda} =_{\beta\eta} uv$$
by the induction hypothesis.

- If $t = \lambda x.\, x$, then
$$(\lambda x.\, x)_{c\lambda} = \mathbf{I}_\lambda = \lambda x.\, x.$$

- If $t = \lambda x.\, y$, then
$$(\lambda x.\, y)_{c\lambda} = (\mathbf{K}y)_\lambda = \mathbf{K}_\lambda y_\lambda = (\lambda zx.\, z)y =_\beta \lambda x.\, y.$$

- If $t = \lambda x.\, uv$, then

$$\begin{array}{rll}(\lambda x.\, uv)_{c\lambda} z & = & (\mathbf{S}(\lambda x.\, u)_c(\lambda x.\, v)_c)_\lambda z \\ & = & \mathbf{S}_\lambda(\lambda x.\, u)_{c\lambda}(\lambda x.\, v)_{c\lambda} z \\ & =_{\beta\eta} & \mathbf{S}_\lambda(\lambda x.\, u)(\lambda x.\, v) z \qquad \text{by the induction hypothesis} \\ & =_\beta & ((\lambda x.\, u)z)((\lambda x.\, v)z) \\ & = & (u[x := z])(v[x := z]) \\ & = & (uv)[x := z] \\ & =_\beta & (\lambda x.\, uv)z.\end{array}$$

  By extensionality, $(\lambda x.\, uv)_{c\lambda} =_{\beta\eta} \lambda x.\, uv$.

- If $t = \lambda x.\, u$ and u is a λ-abstraction, then

$$(\lambda x.\, u)_{c\lambda} = (\lambda x.\, u_c)_{c\lambda} =_{\beta\eta} \lambda x.\, u_{c\lambda} =_{\beta\eta} \lambda x.\, u$$

  by 8.6.8.2 and the induction hypothesis. □

Normal Forms

We do not repeat here the whole story about existence and uniqueness of normal forms told in Sections 3–5, and content ourselves with extending to $\beta\eta$-reducibility the strongest results there obtained for β-reducibility.

Theorem 8.6.11 Strong $\beta\eta$-Normalization. *Every typed λ-term is strongly $\beta\eta$-normalizable.*

Proof. We define the class of computable terms as in 8.3.4. The proofs of propositions 8.3.5, 8.3.7 and 8.3.8 remain valid, and we only need to supplement the proof of 8.3.6 as follows.

In the last paragraph of the proof, we must allow for the additional possibility that a reduction in $(\lambda x^\alpha.\, u^\beta)v^\alpha u_1^{\alpha_1} \cdots u_n^{\alpha_n}$ reduces u^β to a term of the form $t^{\alpha\to\beta}x^\alpha$, with x^α not free in $t^{\alpha\to\beta}$, so that $(\lambda x^\alpha.\, t^{\alpha\to\beta}x^\alpha)v^\alpha u_1^{\alpha_1} \cdots u_n^{\alpha_n}$ reduces to $t^{\alpha\to\beta}v^\alpha u_1^{\alpha_1} \cdots u_n^{\alpha_n}$. However, since x^α is not free in $t^{\alpha\to\beta}$, the latter term is equal to $(t^{\alpha\to\beta}x^\alpha)[x^\alpha := v^\alpha]u_1^{\alpha_1} \cdots u_n^{\alpha_n}$, i.e. any reduction from it can also be performed inside $u^\beta[x^\alpha := v^\alpha]u_1^{\alpha_1} \cdots u_n^{\alpha_n}$. Thus the η-rule does not introduce additional complications. □

Having thus proved the existence of normal forms, we now turn to their uniqueness. We first extend the notion of parallel β-reducibility $\Longrightarrow_\beta$ as we did for the notion of β-reducibility.

Definition 8.6.12 Parallel $\beta\eta$-Reducibility (Tait, Martin-Löf) *The reducibility $\Longrightarrow_{\beta\eta}$ is defined inductively, by adding to 8.5.2 the following clause:*

$$\frac{u \Longrightarrow_{\beta\eta} u_1}{\lambda x.\, ux \Longrightarrow_{\beta\eta} u_1,} \tag{8.19}$$

when x is not free in u.

As a special case of equation 8.19 we have

$$\frac{u \Longrightarrow_{\beta\eta} u}{\lambda x.\, ux \Longrightarrow_{\beta\eta} u,}$$

where the top line consists of an axiom. Thus $\longrightarrow_{1\eta}$ implies $\Longrightarrow_{\beta\eta}$. Inductively, we easily see that $\Longrightarrow_{\beta\eta}$ implies $\longrightarrow_{\beta\eta}$. Then $\longrightarrow_{\beta\eta}$ *is the transitive closure of* $\Longrightarrow_{\beta\eta}$, since it is the transitive closure of $\longrightarrow_{1\beta\eta}$.

Proposition 8.6.13 (Tait, Martin-Löf) *The following is a derived rule for $\Longrightarrow_{\beta\eta}$:*

$$\frac{u \Longrightarrow_{\beta\eta} u^* \qquad t \Longrightarrow_{\beta\eta} t^*}{u[x := t] \Longrightarrow_{\beta\eta} u^*[x := t^*].}$$

Proof. We only need to supplement the proof of 8.5.3 by the following case:

5. $\lambda y.\, u_1 y \Longrightarrow_{\beta\eta} u_2$, with $u_1 \Longrightarrow_{\beta\eta} u_2$
   By definition of substitution:

$$\begin{array}{rcccl} u[x := t] & = & (\lambda y.\, u_1 y)[x := t] & = & \lambda y.\, (u_1[x := t])y \\ u^*[x := t^*] & = & u_2[x := t^*]. & & \end{array}$$

And

$$\cfrac{\cfrac{u_1 \Longrightarrow_{\beta\eta} u_2 \qquad t \Longrightarrow_{\beta\eta} t^*}{u_1[x := t] \Longrightarrow_{\beta\eta} u_2[x := t^*]}}{\lambda y.\, (u_1[x := t])y \Longrightarrow_{\beta\eta} u_2[x := t^*]}$$

by hypothesis, induction hypothesis, and 8.19 of 8.6.12. □

Proposition 8.6.14 Strong Diamond Property for $\Longrightarrow_{\beta\eta}$ (Takahashi [1995]) *For any term t there is a term t^* such that if $t \Longrightarrow_{\beta\eta} t_1$, then $t_1 \Longrightarrow_{\beta\eta} t^*$.*

Proof. We only need to supplement the proof of 8.5.4 by the following definition:

$$t^* = \begin{cases} \lambda x.\, u^* & \text{if } t = \lambda x.\, u \text{ with } u \text{ not of the form } u_1 x, \text{ with } x \text{ not free in } u_1 \\ u^* & \text{if } t = \lambda x.\, ux \text{ and } x \text{ not free in } u \end{cases}$$

and the following case:

5. $t = \lambda x.\, ux$, with x not free in u
   If $t \Longrightarrow_{\beta\eta} t_1$, there are three subcases:

   (a) $t_1 = \lambda x.\, u_1 x$, with $u \Longrightarrow_{\beta\eta} u_1$
   Then

   $$\frac{u_1 \Longrightarrow_{\beta\eta} u^*}{t_1 = \lambda x.\, u_1 x \Longrightarrow_{\beta\eta} u^* = t^*}$$

   by the induction hypothesis and 8.19 of 8.6.12.

   (b) $t_1 = u_1$, with $u \Longrightarrow_{\beta\eta} u_1$
   Then

   $$t_1 = u_1 \Longrightarrow_{\beta\eta} u^* = t^*$$

   by the induction hypothesis.

   (c) $u = \lambda y.\, v$ and $t_1 = \lambda x.\, v_1[y := x] = \lambda y.\, v_1$, with $v \Longrightarrow_{\beta\eta} v_1$
   Then

   $$\frac{v \Longrightarrow_{\beta\eta} v_1}{\lambda y.\, v \Longrightarrow_{\beta\eta} \lambda y.\, v_1}$$

   by 8.13 of 8.5.2, and hence

   $$t_1 = \lambda y.\, v_1 \Longrightarrow_{\beta\eta} (\lambda y.\, v)^* = u^* = t^*$$

   by the induction hypothesis. □

Theorem 8.6.15 Diamond Property for $\beta\eta$-Reduction (Curry and Feys [1958]) *If t_1 and t_2 are terms obtained from t by $\longrightarrow_{\beta\eta}$, then there is a term t^* which can be obtained from t_1 and t_2 by $\longrightarrow_{\beta\eta}$.*

Corollary 8.6.16 Uniqueness of $\beta\eta$-Normal Forms. *Every typed λ-term has exactly one $\beta\eta$-normal form.*

The proof of the Diamond Property given above is useful because it works unchanged for the Untyped Lambda Calculus as well. However, the next exercise shows that for the Typed Lambda Calculus alone a much simpler proof is possible.

Exercise 8.6.17 *The Diamond Property for $\beta\eta$-reduction can be proved as in the proof of Strong Normalization.* (Statman [1985]) (Hint: define a class $\mathcal{C}$ of terms as in 8.3.4, by changing the definition at atomic types α to:

$$t^\alpha \in \mathcal{C} \;\Leftrightarrow\; \text{the Diamond Property holds when starting from } t^\alpha.$$

The definition at arrow types is the same. We then prove by induction on types that for every term t in $\mathcal{C}$ the Diamond Property holds when starting from t, and by induction on terms that every typed term is in $\mathcal{C}$. For the first part, we need to reduce the type of $t^{\alpha\to\beta}$. By the induction hypothesis, the Diamond Property holds when starting from $t^{\alpha\to\beta}x^\alpha$. By inserting λ's everywhere, the Diamond Property holds when starting from $\lambda x^\alpha.\, t^{\alpha\to\beta}x^\alpha$. By the η-rule, the latter is equal to $t^{\alpha\to\beta}$ itself.)

æ

Chapter 9

The Curry-Howard Isomorphism

In the previuous chapters we introduced the Typed Lambda Calculus independently of logic, as a theory of computable functions. We now look at the relationships.

Such a relationship is beneficial in both directions: on the one hand logic can be seen in a computational way (as a calculus whose objects are proofs), and on the other hand the simple system of types of Chapter C can be extended to reflect various well-known systems of logic, in a process that will lead us to powerful systems.

THREE THEORIES:

1. arrow types: implication

2. arrow and product types: implication and conjunction

3. surjective pairs: symmetric reduction rules

9.1 The Curry-Howard Isomorphism

Theorem 9.1.1 The Curry-Howard Isomorphism (??) *There is an isomorphism between:*

- *the intuitionistic proof theory of implication and conjunction, with the Natural Deduction rules of introduction, elimination, normalization and inverse normalization;*

- *the reduction theory of the Typed Lambda Calculus, with the β-rule and surjective pairs.*

To know whether a formula has a proof: look for a term with the formula as type.

To know whether a type admits a terms with that type (is inhabited) see whether the type is provable (as a formula).

9.2 Normal forms

Strong Normalization and uniqueness of normal form for proof in $\mathcal{N}$.

Discussion of the relationships with Cut Elimination: non uniqueness.

9.3 Semantics

Turning models of the Typed Lambda Calculus into models of Implicational Logic and conversely.

Proving strong normalization by Kripke Models (Gallier).

9.4 Complexity

Complexity of reduction procedures and of type inhabitation.

æ

Chapter 10

Semantics

Until now we have developed the Typed Lambda Calculus as a purely syntactical theory, with terms as objects that can be manipulated according to the reduction rules. Intuitively, however, we always had an intended interpretation in mind: terms of nonatomic types were thought of as naming functions, and the β-rule was intended to codify an atomic step in the evaluation of a function at a given argument.

The question of the relationship between the two levels, i.e. the syntactical manipulation of terms and the functional intended meaning, was already indirectly addressed in Chapter 8, where we showed that the Typed Lambda Calculus can be seen as a theory of functions in the abstract sense that the evaluation of a term always terminates, with a uniquely determined answer.

We now address the problem of the relationship between the two levels in a more direct and conventional way, by looking for possible interpretations of typed λ-terms as functions in the usual mathematical sense. This will be achieved by exhibiting a number of possible models.

10.1 Models

In our example of polynomial expressions, we had in mind natural numbers as intended meanings. This can be reformulated by saying that the set N of natural numbers was the 'intended model' for the calculus of polynomial expressions. Of course, it is not enough to specify which objects we had in mind: we also have to say how we associate to every polynomial expression a number that such an expression describes.

This is quite clear if the expression has no variables: then the intented meaning is the *number* described by that expression, e.g. 44 in the case of $3^2 + 5 \cdot 7$. If there are variables, the situation is slightly more complicated: the meaning will ultimately depend on the value assigned to the variables, e.g. $x^2 + 5 \cdot 7$ will denote a number for any fixed x, and hence by itself it denotes a generic number of a certain form, i.e. a *function* from N to N.

An obvious property of the 'meaning' of polynomial expressions is that they coincide for polynomial expressions that can be obtained one from the other by formal manipulations, i.e. for polynomial expressions with the same normal form.

All this can now be generalized to λ-terms, in a natural way. First, we have here terms of different possible types, and thus we need a set M_α of 'meanings' for each type α. We further need to interpret terms t^α as descriptions of elements $[\![t^\alpha]\!]$ of M_α. Recall that, from the start, terms have been taken as descriptions of objects. The novelty here is that we now specify which objects we have in mind (namely, elements of M_α).

The fact that a description t^α could be either a generic or a proper name, depending on whether t^α has free variables or not, is reflected in the fact that its interpretation $[\![t^\alpha]\!]$ is either a generic object of M_α, i.e. the value of a function with values in M_α, or a specific object of M_α. More precisely, if t^α is closed, then $[\![t^\alpha]\!] \in M_\alpha$. And if t^α has free variables $x_1^{\alpha_1}, \ldots, x_n^{\alpha_n}$, then $[\![t^\alpha]\!]$ is a function

$$[\![t^\alpha]\!] : M_{\alpha_1} \times \cdots \times M_{\alpha_n} \longrightarrow M_\alpha.$$

Since λ-abstraction allows us to abstract x^α in u^β even if the former does not occur free in the latter, we will need to consider $[\![t^\alpha]\!]$ as such a function also when the free variables of t^α are only *among* $x_1^{\alpha_1}, \ldots, x_n^{\alpha_n}$. This is somewhat unsatisfactory, since the same interpretation is considered as a different function in different situations.

Finally, the interpretation function should respect β-equality, i.e. we should have

$$t_1^\alpha =_\beta t_2^\alpha \;\Rightarrow\; [\![t_1^\alpha]\!] = [\![t_2^\alpha]\!],$$

where the equality on the right indicates identity of mathematical objects. This is somewhat unclear, since e.g. t_1 might contain free variables while t_2 might be closed: in the first case the interpretation will be an element, in the second a function, and we should then explain how they could be considered 'equal'.

Environments

The treatment of free variables, which is the source of some of the unsatisfactory features of the formulation above, is much simplified by a simple observation. Instead of considering the interpretation of a term as a function of *some* variables, we consider it as a function of *all* (infinitely many!) variables. This is possible because, instead of individual variables ranging separately over the appropriate domains, we can consider the set of all variables ranging globally over all possible domains. This process is easily described in terms of *environments* for $\{M_\alpha\}_\alpha$, i.e. interpretations of every variable by an element of $\{M_\alpha\}_\alpha$ of the appropriate level.

Definition 10.1.1 *An* **environment** *for $\{M_\alpha\}_\alpha$ is a function*

$$\rho : \textit{Variables} \longrightarrow \bigcup_\alpha M_\alpha$$

such that

$$\rho(x^\alpha) \in M_\alpha.$$

We denote by $\boldsymbol{\mathcal{E}}$ *the set of environments, by* $\boldsymbol{\rho}$ *an environment, and by* $\boldsymbol{\rho[x^\alpha := a]}$ *the modification of the environment ρ obtained by interpreting x^α as $a \in M_\alpha$, i.e.*

$$(\rho[x^\alpha := a])(y^\beta) = \begin{cases} a & \textit{if } y^\beta = x^\alpha \\ \rho(y^\beta) & \textit{otherwise.} \end{cases}$$

Notice that the two notations

$$u[x := v] \qquad \text{and} \qquad \rho[x := a]$$

are related, as the similarity of notation underlines, but different. The first denotes a syntactical substitution of the term v for the variable x in the term u. The second denotes a semantical assignment of the value a to the variable x in the environment function ρ.

The advantage of environments is simply *uniformity of notation*. They allow us to consider the interpretation of a term t^α not as a function

$$[\![t^\alpha]\!] : M_{\alpha_1} \times \cdots \times M_{\alpha_n} \longrightarrow M_\alpha,$$

but rather as a function of all variables

$$[\![t^\alpha]\!] : \mathcal{E} \longrightarrow M_\alpha.$$

We will write $[\![t^\alpha]\!]_\rho$ for the interpretation of t^α under the environment ρ.

Models

We can now use environments to restate the previous description of interpretation in a more compact and transparent way.

Definition 10.1.2 *A* **model** *of the Typed Lambda Calculus is a structure*

$$\mathcal{M} = \langle \{M_\alpha\}_\alpha, [\![\ \ ]\!]^{\mathcal{M}} \rangle$$

with the following properties, where ρ is any environment for $\{M_\alpha\}_\alpha$:

1. *if t^α is a term of type α, then $(\forall \rho)([\![t^\alpha]\!]^{\mathcal{M}}_\rho \in M_\alpha)$.*

2. $[\![\ ]\!]^{\mathcal{M}}$ *respects* β*-equality, i.e.*

$$t_1^\alpha =_\beta t_2^\alpha \;\Rightarrow\; (\forall\rho)([\![t_1^\alpha]\!]_\rho^{\mathcal{M}} = [\![t_2^\alpha]\!]_\rho^{\mathcal{M}}),$$

where the equality on the right indicates identity of objects in M_α*.*

An **extensional model** *is a model that respects* $\beta\eta$*-equality, i.e.*

$$t_1^\alpha =_{\beta\eta} t_2^\alpha \;\Rightarrow\; (\forall\rho)([\![t_1^\alpha]\!]_\rho^{\mathcal{M}} = [\![t_2^\alpha]\!]_\rho^{\mathcal{M}}).$$

To improve readability we can omit either of the indeces ρ or $\mathcal{M}$ in $[\![\ ]\!]_\rho^{\mathcal{M}}$, when no confusion arises. Moreover, we will also write

$$\mathcal{M} \models t_1^\alpha = t_2^\alpha \qquad \text{for} \qquad (\forall\rho)([\![t_1^\alpha]\!]_\rho^{\mathcal{M}} = [\![t_2^\alpha]\!]_\rho^{\mathcal{M}}),$$

so that the soundness conditions in the definition of a model can be stated more succinctly as

$$t_1^\alpha =_\beta t_2^\alpha \;\Rightarrow\; \mathcal{M} \models t_1^\alpha = t_2^\alpha$$

or

$$t_1^\alpha =_{\beta\eta} t_2^\alpha \;\Rightarrow\; \mathcal{M} \models t_1^\alpha = t_2^\alpha.$$

10.2 Term Models

The example of polynomial expressions shows how a model can easily be obtained by identifying the objects of the calculus with their normal forms. Thus, for example, we can write both

$$[\![3^2 + 5 \cdot 7]\!]^N = [\![44]\!]^N,$$

to mean that $3^2 + 5 \cdot 7$ and 44 are descriptions of the same natural number, and

$$[\![3^2 + 5 \cdot 7]\!]^N = 44,$$

to mean that $3^2 + 5 \cdot 7$ is a description of the number 44. This corresponds to the discussion above, on how numbers can be taken not only to be *described* by normal expressions in decimal notation, but to *be* such expressions.

This can be done for the Typed Lambda Calculus as well, by taking terms in normal form to be not only as particularly simple descriptions of objects, but as the objects themselves. The interpretation of a term will thus be its normal form.

Definition 10.2.1 *The* **first term model** $\mathcal{T}_1$ *is defined as follows:*

1. The underlying structure consists of:

$$T_\alpha = \{\textit{terms of type } \alpha \textit{ in normal form}\}$$

2. *Let ρ be an environment for $\{T_\alpha\}_\alpha$, i.e. a function assigning to every variable a term in normal form of the same type. Then*

$$[\![t^\alpha]\!]^{\mathcal{T}_1}_\rho = \textit{the normal form of } t^\alpha[\vec{x} := \rho(\vec{x})],$$

where $t^\alpha[\vec{x} := \rho(\vec{x})]$ indicates the result of the simultaneous substitution of the term $\rho(x)$ for any free variable x of t^α.

Notice that it is not enough to consider only the *closed* terms, because in a pure theory there is no closed term of atomic type.

The next result shows not only that $\mathcal{T}_1$ is a model of λ-calculus, i.e. β-equal terms have the same interpretation under every assignment, but also that β-different terms have different interpretations under at least one assignment. This property of models is related to what is called **full abstraction** in the literature.

Proposition 10.2.2 *The structure $\mathcal{T}_1$ is a model of the Typed Lambda Calculus. Actually,*

$$t_1^\alpha =_\beta t_2^\alpha \Leftrightarrow \mathcal{T}_1 \models t_1^\alpha = t_2^\alpha.$$

Proof. If $t_1^\alpha =_\beta t_2^\alpha$, then

$$t_1^\alpha[\vec{x} := \rho(\vec{x})] =_\beta t_2^\alpha[\vec{x} := \rho(\vec{x})],$$

because $=_\beta$ is invariant under simultaneous substitutions. But β-equal terms have the same normal form, and thus

$$[\![t_1^\alpha]\!]_\rho = [\![t_2^\alpha]\!]_\rho.$$

Conversely, suppose $t_1^\alpha \neq_\beta t_2^\alpha$. Then they have different normal forms. If ρ is the identity function, which is an environment because the variables are terms in normal form, then

$$t_1^\alpha[\vec{x} := \rho(\vec{x})] = t_1^\alpha \neq_\beta t_2^\alpha = t_2^\alpha[\vec{x} := \rho(\vec{x})].$$

Thus

$$[\![t_1^\alpha]\!]_\rho \neq [\![t_2^\alpha]\!]_\rho,$$

because $t_1^\alpha[\vec{x} := \rho(\vec{x})]$ and $t_2^\alpha[\vec{x} := \rho(\vec{x})]$ have different normal forms. □

The first term model is very natural, but to interpret a term as a normal form obviously requires the existence of normal forms. Thus the model will have no parallel in the Untyped Lambda Calculus. However, a simple modification of it will. Since a model has to identify terms that are β-equal, the new idea is to consider the set of all terms that are β-equal to a given term as an interpretation of it. Technically, we consider equivalence classes of terms w.r.t. the equivalence relation $=_\beta$.

Definition 10.2.3 *The* **second term model** $\mathcal{T}_2$ *is defined as follows:*

1. *The underlying structure consists of:*

$$T_\alpha = \{\textit{equivalence classes of terms of type } \alpha \textit{ w.r.t. } =_\beta\}$$

2. *Given an environment ρ on $\{T_\alpha\}_\alpha$, i.e. a function assigning to every variable the equivalence class of a term of the same type, let ρ^* be a choice function for ρ, i.e. a function that associates to every variable x a term in the equivalence class $\rho(x)$. Then*

$$[\![t^\alpha]\!]_\rho^{\mathcal{T}_2} = \textit{the equivalence class of } t^\alpha[\vec{x} := \rho^*(\vec{x})],$$

where $t^\alpha[\vec{x} := \rho^*(\vec{x})]$ indicates the result of the simultaneous substitution of the term $\rho^*(x)$ for any free variable x of t^α.

Proposition 10.2.4 *The structure $\mathcal{T}_2$ is a model of the Typed Lambda Calculus. Actually,*

$$t_1^\alpha =_\beta t_2^\alpha \;\Leftrightarrow\; \mathcal{T}_2 \models t_1^\alpha = t_2^\alpha.$$

Proof. If $t_1^\alpha =_\beta t_2^\alpha$, then

$$t_1^\alpha[\vec{x} := \rho^*(\vec{x})] =_\beta t_2^\alpha[\vec{x} := \rho^*(\vec{x})],$$

because $=_\beta$ is invariant under simultaneous substitutions. But β-equal terms are in the same equivalence class, and thus

$$[\![t_1^\alpha]\!]_\rho = [\![t_2^\alpha]\!]_\rho.$$

Conversely, suppose $t_1^\alpha \neq_\beta t_2^\alpha$. Then they are in different equivalence classes. If ρ is the environment associating its equivalence class to every variable, we can choose as ρ^* the identity function. Then

$$t_1^\alpha[\vec{x} := \rho^*(\vec{x})] = t_1^\alpha \neq_\beta t_2^\alpha = t_2^\alpha[\vec{x} := \rho^*(\vec{x})],$$

and hence

$$[\![t_1^\alpha]\!]_\rho \neq [\![t_2^\alpha]\!]_\rho,$$

because $t_1^\alpha[\vec{x} := \rho^*(\vec{x})]$ and $t_2^\alpha[\vec{x} := \rho^*(\vec{x})]$ are in different equivalence classes. □

Extensional term models can be defined similarly, by considering $\beta\eta$ normal forms and $\beta\eta$-equality, respectively.

Despite their advantages, the term models provide semantical interpretations of terms too close to the original syntactical presentation. In other words, they are well-behaved but not very insightful. This is the reason to continue, in the next sections, the search for other more informative models.

10.3 Functional Models

We now turn to the consideration of models in which terms of arrow types are interpreted as functions in the usual mathematical sense, and terms of atomic types are interpreted as elements of given sets. In particular, $M_{\alpha\to\beta}$ is here considered as a subset of the set $(M_\beta)^{M_\alpha}$ of all functions from M_α to M_β.

In practice, we define the interpretation function $[\![\ ]\!]^{\mathcal{M}}$ in the following canonical way, which reduces it to the *definition* of a structure $\{M_\alpha\}_\alpha$, and the *verification* of closure under informal abstraction.

Definition 10.3.1 Canonical Interpretation. *Given $\{M_\alpha\}_\alpha$ such that, for every α and β,*

$$M_{\alpha\to\beta} \subseteq (M_\beta)^{M_\alpha},$$

and an environment ρ on it, we define $[\![\ ]\!]_\rho$ by induction on terms, as follows:

$$[\![t^\alpha]\!]_\rho = \begin{cases} \rho(x^\alpha) & \textit{if } t^\alpha = x^\alpha \\ [\![u^{\gamma\to\alpha}]\!]_\rho([\![v^\gamma]\!]_\rho) & \textit{if } t^\alpha = u^{\gamma\to\alpha}v^\gamma \\ \Lambda X^\gamma.\, [\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]} & \textit{if } t^\alpha = \lambda x^\gamma.\, u^\beta, \end{cases}$$

where $\Lambda X^\gamma.\, [\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]}$ denotes the function

$$a \in M_\gamma \longmapsto [\![u^\beta]\!]_{\rho[x^\gamma:=a]} \in M_\beta.$$

Note that Λ is simply shorthand for 'the function with argument ... and value ...'. The use of a different symbol stresses the fact that Λ is an informal abstraction operator that produces names for mathematical *functions* on the structure $\{M_\alpha\}_\alpha$, while λ is a formal operator that produces names for *terms*. Obviously the two uses are related, in the sense that λ is meant to formally capture some intuitive properties of Λ.

A similar discussion holds for X^γ, which is used as an informal variable ranging over *elements* of M_γ, while x^γ is a formal variable ranging over *terms* of type γ. Again the two uses are related, in the sense that elements of M_γ are meant to interpret terms of type γ.

We should check, inductively, that $[\![t^\alpha]\!]_\rho$ is a member of M_α. In the first clause, it is so by definition of environment. In the second clause, it is so by the induction hypothesis, since then $[\![u^{\gamma\to\alpha}]\!]_\rho$ is an element of $M_{\gamma\to\alpha}$, and hence a function from M_γ to M_α. In the last case, by the induction hypothesis, we certainly obtain a function from M_γ to M_β, but not necessarily an element of $M_{\gamma\to\beta}$, without further hypotheses on $\{M_\alpha\}_\alpha$. For specific structures, this will have to be *verified*.

Full functional models

The simplest (and most simpleminded) class of functional models we can think of is obtained by not imposing any restriction on the class of functions we consider at arrow type levels. The underlying structure of such a model is defined as follows.

Definition 10.3.2 *For any set A, the* **full type hierarchy over $\boldsymbol{A}$** *is the structure $\{A_\alpha\}_\alpha$ defined as follows:*

1. *$A_\alpha = A$ for α atomic*
2. *$A_{\alpha\to\beta}$ is the full function space, i.e. the set $A_\beta^{A_\alpha}$ of all functions from A_α to A_β.*

The definition of the model becomes then the following.

Definition 10.3.3 *The* **full functional model over $\boldsymbol{A}$** *is the structure*

$$\mathcal{F}_A = \langle \{A_\alpha\}_\alpha, [\![\ \ ]\!]^{\mathcal{F}_A} \rangle,$$

where $[\![\ \ ]\!]^{\mathcal{F}_A}$ is the canonical interpretation over $\{A_\alpha\}_\alpha$.

The main point here is that the interpretation of $[\![\lambda x^\gamma.\, u^\beta]\!]_\rho$, which by definition is a function from A_γ to A_β, is now an element of $A_{\gamma\to\beta}$ because the latter contains *all* such functions. Thus $[\![\ \ ]\!]^{\mathcal{F}_A}$ is automatically well-defined.

Theorem 10.3.4 Soundness for Full Functional Models. *For any nonempty set A, the structure $\mathcal{F}_A$ is an extensional model of the Typed Lambda Calculus. More precisely,*

$$t_1^\alpha =_{\beta\eta} t_2^\alpha \ \Rightarrow\ \mathcal{F}_A \models t_1^\alpha = t_2^\alpha.$$

Proof. Since two terms are $\beta\eta$-equal if they reduce to the same term, it is enough to prove the result when $t_1^\alpha \longrightarrow_{\beta\eta} t_2^\alpha$. By induction on the number of steps, it is enough to prove the result when $t_1^\alpha \longrightarrow_{1\beta\eta} t_2^\alpha$. This is proved by induction on the definition of $\longrightarrow_{1\beta\eta}$. There are five cases, the first two of which are the crucial verifications, while the three remaining ones hold trivially.

1. $(\lambda x^\alpha.\, u^\beta)v^\alpha \longrightarrow_{1\beta\eta} u^\beta[x^\alpha := v^\alpha]$
   Then

$$\begin{aligned}
[\![(\lambda x^\alpha.\, u^\beta)v^\alpha]\!]_\rho &= [\![\lambda x^\alpha.\, u^\beta]\!]_\rho([\![v^\alpha]\!]_\rho) \\
&= (\Lambda X^\alpha.\, [\![u^\beta]\!]_{\rho[x^\alpha := X^\alpha]})([\![v^\alpha]\!]_\rho) \\
&= [\![u^\beta]\!]_{\rho[x^\alpha := [\![v^\alpha]\!]_\rho]} \\
&= [\![u^\beta[x^\alpha := v^\alpha]]\!]_\rho.
\end{aligned}$$

The first two equalities hold by definition of $[\![\ ]\!]$. The third holds by definition of Λ as a mathematical function, which is computed by instantiating the variable to the argument. The last equality holds by the fact, proved in 10.3.5, that substitutions in terms and in environments commute.

2. $\lambda x^\alpha.\, t^{\alpha\to\beta} x^\alpha \longrightarrow_{1\beta\eta} t^{\alpha\to\beta}$
   Then

$$\begin{array}{rcl} [\![\lambda x^\alpha.\, t^{\alpha\to\beta} x^\alpha]\!]_\rho & = & \Lambda X^\alpha.\, [\![t^{\alpha\to\beta} x^\alpha]\!]_{\rho[x^\alpha := X^\alpha]} \\ & = & \Lambda X^\alpha.\, [\![t^{\alpha\to\beta}]\!]_{\rho[x^\alpha := X^\alpha]}([\![x^\alpha]\!]_{\rho[x^\alpha := X^\alpha]}) \\ & = & \Lambda X^\alpha.\, [\![t^{\alpha\to\beta}]\!]_{\rho[x^\alpha := X^\alpha]}(X^\alpha) \\ & = & [\![t^{\alpha\to\beta}]\!]_\rho. \end{array}$$

   The first three equalities hold by definition of $[\![\ ]\!]$. The last holds by definition of Λ as a mathematical function.

3. $u_1^{\alpha\to\beta} v^\alpha \longrightarrow_{1\beta\eta} u_2^{\alpha\to\beta} v^\alpha$
   By the induction hypothesis, $u_1^{\alpha\to\beta\eta} \longrightarrow_{1\beta} u_2^{\alpha\to\beta}$. Hence, for any ρ,

$$[\![u_1]\!]_\rho = [\![u_2]\!]_\rho.$$

   Then the two interpretations are the same function, and have the same value for the same argument $[\![v]\!]_\rho$. Hence

$$[\![u_1 v]\!]_\rho = [\![u_1]\!]_\rho([\![v]\!]_\rho) = [\![u_2]\!]_\rho([\![v]\!]_\rho) = [\![u_2 v]\!]_\rho,$$

   where the two outer equalities hold by definition of $[\![\ ]\!]$.

4. $u^{\alpha\to\beta} v_1^\alpha \longrightarrow_{1\beta\eta} u^{\alpha\to\beta} v_2^\alpha$
   By the induction hypothesis, $v_1^\alpha \longrightarrow_{1\beta\eta} v_2^\alpha$. Hence, for any ρ,

$$[\![v_1]\!]_\rho = [\![v_2]\!]_\rho.$$

   Then the two interpretations are the same element, and any function $[\![u]\!]_\rho$ applied to them has the same value. Hence

$$[\![u v_1]\!]_\rho = [\![u]\!]_\rho([\![v_1]\!]_\rho) = [\![u]\!]_\rho([\![v_2]\!]_\rho) = [\![u v_2]\!]_\rho,$$

   where the two outer equalities hold by definition of $[\![\ ]\!]$.

5. $\lambda x^\alpha.\, u_1^\beta \longrightarrow_{1\beta\eta} \lambda x^\alpha.\, u_2^\beta$
   By the induction hypothesis, $u_1^\beta \longrightarrow_{1\beta\eta} u_2^\beta$. Hence, for any ρ,

$$[\![u_1]\!]_\rho = [\![u_2]\!]_\rho.$$

In particular, for any $a^\alpha \in A_\alpha$,

$$[\![u_1]\!]_{\rho[x^\alpha := a^\alpha]} = [\![u_2]\!]_{\rho[x^\alpha := a^\alpha]}.$$

Then

$$[\![\lambda x^\alpha.\, u_1^\beta]\!] = \Lambda X^\alpha.\, [\![u_1^\beta]\!]_{\rho[x^\alpha := X^\alpha]} = \Lambda X^\alpha.\, [\![u_2^\beta]\!]_{\rho[x^\alpha := X^\alpha]} = [\![\lambda x^\alpha.\, u_2^\beta]\!],$$

where the two outer equalities hold by definition of $[\![\ ]\!]$. □

Exercise 10.3.5 Substitution Lemma. *For any u, v and ρ,*

$$[\![u]\!]_{\rho[x := [\![v]\!]_\rho]} = [\![u[x := v]]\!]_\rho.$$

(Hint: by induction on u. The only non trivial case is $u = \lambda y.\, u_1$. By possibly using the α-rule, we may suppose that y is not free in v and is distinct from x. Then

$$\begin{array}{rcl} [\![u[x := v]]\!] & = & \Lambda Y.\, [\![u_1[x := v]]\!]_{\rho[y := Y]} \\ & = & \Lambda Y.\, [\![u_1]\!]_{(\rho[y:=Y])[x := [\![v]\!]_{\rho[y:=Y]}]} \\ & = & \Lambda Y.\, [\![u_1]\!]_{(\rho[y:=Y])[x := [\![v]\!]_\rho]} \\ & = & \Lambda Y.\, [\![u_1]\!]_{(\rho[x := [\![v]\!]_\rho])[y:=Y]} \\ & = & [\![u]\!]_{\rho[x := [\![v]\!]_\rho]}. \end{array}$$

The first equality holds by definition of $[\![\ ]\!]_\rho$, because $u[x := v] = \lambda y.\, u_1[x := v]$). The second holds by the induction hypothesis applied to the environment $\rho[y := Y]$. The third holds because y does not occur free in v. The fourth holds because y and x are distinct. The last holds by definition of $[\![\ ]\!]_{\rho[x := [\![v]\!]_\rho]}$.)

Completeness

If A is a *finite* set, we do not expect the opposite implication to hold in the previous result, i.e. $\mathcal{F}_A$ to be able to distinguish among different terms. Indeed, if A is finite, then $A_{(p \to p) \to (p \to p)}$ is finite, but there are infinitely many distinct closed terms of type $(p \to p) \to (p \to p)$. For example, the following:

$$\lambda f^{p \to p} x^p.\, \underbrace{f(\cdots(f}_{n \text{ times}}(x))\cdots).$$

However, the next result proves that the finiteness of A is the only obstacle.

Theorem 10.3.6 Completeness for Full Functional Models (Friedman [1975]) *For any infinite set A,*

$$t_1^\alpha =_{\beta\eta} t_2^\alpha \;\Leftrightarrow\; \mathcal{F}_A \models t_1^\alpha = t_2^\alpha.$$

Proof. The left to right direction has already been proved in 10.3.4.

For the right to left direction, we take advantage of the fact that two terms are $\beta\eta$-equal if and only if they have the same normal form, and rewrite the requirement

$$\mathcal{F}_A \models t_1^\alpha = t_2^\alpha \;\Rightarrow\; t_1^\alpha =_{\beta\eta} t_2^\alpha$$

as

$$\mathcal{F}_A \models t_1^\alpha = t_2^\alpha \;\Rightarrow\; t_1^\alpha \text{ and } t_2^\alpha \text{ have the same normal form.}$$

To satisfy this, it is enough to find a *single* environment ρ on $\mathcal{F}_A$ such that

$$[\![t_1^\alpha]\!]_\rho = [\![t_2^\alpha]\!]_\rho \;\Rightarrow\; t_1^\alpha \text{ and } t_2^\alpha \text{ have the same normal form.}$$

We are thus looking for an environment ρ and a function I such that

$$I([\![t^\alpha]\!]_\rho) = \text{the normal form of } t^\alpha.$$

From this the result follows, since

$$\begin{array}{rcl} \mathcal{F}_A \models t_1^\alpha = t_2^\alpha & \Rightarrow & [\![t_1^\alpha]\!]_\rho = [\![t_2^\alpha]\!]_\rho \\ & \Rightarrow & I([\![t_1^\alpha]\!]_\rho) = I([\![t_2^\alpha]\!]_\rho) \\ & \Rightarrow & t_1^\alpha \text{ and } t_2^\alpha \text{ have the same normal form} \\ & \Rightarrow & t_1^\alpha =_{\beta\eta} t_2^\alpha. \end{array}$$

Since I is defined on $\bigcup_\alpha A_\alpha$, i.e. on interpretation of terms, it can be seen as a canonical *inversion map for the interpretation function* $[\![\;]\!]_\rho$. Moreover:

- To ensure that I is properly defined on interpretations, we require it to be *level preserving*, and thus to be a family of functions I_α sending elements of A_α to terms in normal form of type α.

- Since not every element of A_α can be the interpretation of a term of type α,[1] I_α will in general be *partial*.

- To ensure that every interpretation is inverted, I_α has to be *onto* the set of terms of type α and in normal form.

- Finally, since the family $\{I_\alpha\}_\alpha$ is intended to invert canonical interpretations, which by definition have the property that

$$[\![u^{\alpha\to\beta}v^\alpha]\!] = [\![u^{\alpha\to\beta}]\!]([\![v^\alpha]\!]),$$

  we require the following *amalgamation property*: if $f \in A_{\alpha\to\beta}$, $a \in A_\alpha$ and the relevant values are defined, then

$$I_\beta(f(a)) =_{\beta\eta} I_{\alpha\to\beta}(f)(I_\alpha(a)).$$

[1] Except possibly for atomic types, each A_α contains uncountably many elements, because A is infinite, while there are only countably many terms.

This concludes the plan of the proof, and we now turn to its implementation.

We start by defining an inversion function $I = \{I_\alpha\}_\alpha$, by induction on types:

- *atomic types* α
  Let I_α be any partial function from A onto the set of terms of type α in normal form. Notice that such an I_α exists because A is infinite. The hypothesis is needed because there are infinitely many terms of type α in normal form, e.g. all distinct variables of type α.

- *arrow types* $\alpha \to \beta$
  If I_α and I_β are partial functions onto the sets of terms in normal form of type α and β, respectively, we define $I_{\alpha\to\beta}$ using the amalgamation property. Given $f \in A_{\alpha\to\beta}$, we let $I_{\alpha\to\beta}(f)$ be any term $t^{\alpha\to\beta}$ in normal form, if it exists, such that

  $$(\forall a \in A_\alpha)[I_\alpha(a) \text{ defined } \Rightarrow I_\beta(f(a)) \simeq \text{the normal form of } tI_\alpha(a)],$$

  where $\simeq$ means 'defined and equal to'.

  First, $I_{\alpha\to\beta}$ is a partial function. Indeed, for any f there is at most one such a term t. Since I_α is onto the set of terms of type α in normal form, the definition of t determines its behaviour on (the normal form of) every term of type α. Any two possible choices of t thus behave the same way on all terms of type α, and they must thus be $\beta\eta$-equal. Since they are in normal form, and normal forms are unique, they must then be the same term.

  Second, $I_{\alpha\to\beta}$ is onto the set of terms of type $\alpha \to \beta$ in normal form. Indeed, for any such term t, we can define $f \in A_{\alpha\to\beta}$ such that $I_{\alpha\to\beta}(f) = t$ as follows. Given $a \in A_\alpha$, if $I_\alpha(a)$ is defined, then it is a term of type α, and $tI_\alpha(a)$ is a term of type β. Since I_β is onto the set of terms of type β in normal form, there is $b_a \in A_\beta$ such that

  $$I_\beta(b_a) = \text{the normal form of } tI_\alpha(a).$$

  Let $f \in A_{\alpha\to\beta}$ be any function that assigns to any $a \in A_\alpha$ the element b_a if $I_\alpha(a)$ is defined, and an arbitrary element of A_β otherwise.

We now define an environment ρ in such a way that I inverts precisely $[\![\ ]\!]_\rho$. Since the interpretation is uniquely determined by the behaviour of ρ on variables, it is enough to define ρ as any assignment such that

$$I_\alpha(\rho(x^\alpha)) = x^\alpha.$$

The definition makes sense because each variable x^α is a term of type α in normal form, and hence it is in the range of I_α. Then ρ simply picks up any element $a \in A_\alpha$ such that $I_\alpha(a) = x^\alpha$.

It remains to check that

$$I_\alpha([\![t^\alpha]\!]_\rho) \simeq \textit{the normal form of } t^\alpha.$$

Since $I_\alpha([\![t^\alpha]\!]_\rho)$ is a term in normal form, when defined, by uniqueness of normal forms it is enough to show that

$$I_\alpha([\![t^\alpha]\!]_\rho) =_{\beta\eta} t^\alpha.$$

We proceed by induction on terms:

1. If $t = x^\alpha$, then
$$I([\![x^\alpha]\!]_\rho) = I(\rho(x^\alpha)) = x^\alpha,$$
where the first equality holds by definition of $[\![\ \ ]\!]_\rho$, and the second by definition of ρ.

2. If $t = u^{\alpha\to\beta}v^\alpha$, then
$$\begin{array}{rcl} I_\beta([\![u^{\alpha\to\beta}v^\alpha]\!]_\rho) & = & I_\beta([\![u^{\alpha\to\beta}]\!]_\rho([\![v^\alpha]\!]_\rho)) \\ & =_{\beta\eta} & I_{\alpha\to\beta}([\![u^{\alpha\to\beta}]\!]_\rho)(I_\alpha([\![v^\alpha]\!]_\rho)) \\ & =_{\beta\eta} & u^{\alpha\to\beta}v^\alpha, \end{array}$$
where the first equality holds by definition of $[\![\ \ ]\!]_\rho$, the second by the amalgamation property, and the last by the induction hypothesis.

3. If $t = \lambda x^\alpha.\, u^\beta$, then we want
$$I_{\alpha\to\beta}([\![\lambda x^\alpha.\, u^\beta]\!]_\rho) =_{\beta\eta} \lambda x^\alpha.\, u^\beta.$$
By extensionality and normalization it is enough to show that, for any term v^α in normal form,
$$I_{\alpha\to\beta}([\![\lambda x^\alpha.\, u^\beta]\!]_\rho)v^\alpha =_{\beta\eta} (\lambda x^\alpha.\, u^\beta)v^\alpha$$
and hence, by the β-rule, that
$$I_{\alpha\to\beta}([\![\lambda x^\alpha.\, u^\beta]\!]_\rho)v^\alpha =_{\beta} u^\beta[x^\alpha := v^\alpha].$$
Since I_α is onto the set of terms of type α in normal form, there is $a \in A_\alpha$ such that $a = [\![v^\alpha]\!]_\rho$ and $I_\alpha(a) = v^\alpha$. Then
$$\begin{array}{rcl} I_{\alpha\to\beta}([\![\lambda x^\alpha.\, u^\beta]\!]_\rho)v^\alpha & = & I_{\alpha\to\beta}([\![\lambda x^\alpha.\, u^\beta]\!]_\rho)I_\alpha(a) \\ & = & I_\beta([\![\lambda x^\alpha.\, u^\beta]\!]_\rho(a)) \\ & = & I_\beta([\![u^\beta]\!]_{\rho[x^\alpha:=a]}) \\ & = & I_\beta([\![u^\beta]\!]_{\rho[x^\alpha:=[\![v^\alpha]\!]_\rho]}) \\ & = & I_\beta([\![u^\beta[x^\alpha := v^\alpha]]\!]_\rho) \\ & =_{\beta\eta} & u^\beta[x^\alpha := v^\alpha], \end{array}$$

where the equalities hold by choice of a, amalgamation property, definition of $[\![\ ]\!]$, choice of a, Substitution Lemma (10.3.5), and the induction hypothesis. □

The previous proof can be abstractly seen as an embedding of the first term model $\mathcal{T}_1$ into a quotient of $\mathcal{F}_A$, with a transfer of the completeness property from the former (for which it holds trivially, see 10.2.2) to the latter. The embedding is provided by the inverse relations I_α^{-1} (since I_α is a partial onto function, its inverse I_α^{-1} is a total one-one relation, but not necessarily a function), and the quotient is generated on each A_α by the partial equivalence relation R_α induced by the partial function I_α as follows: if $a, b \in A_\alpha$, then

$$aR_\alpha b \ \Leftrightarrow\ I_\alpha(a) \simeq I_\alpha(b).$$

The definition of I_α implies that if $f, g \in A_{\alpha\to\beta}$, then

$$fR_{\alpha\to\beta}g \ \Leftrightarrow\ (\forall a, b \in A_\alpha)[aR_\alpha b \ \Rightarrow\ f(a)R_\beta g(b)].$$

Thus, proceeding by induction on types, we consider functions compatible with the equivalence relations previously defined, and identifies all functions that induce the same restriction on the previously defined quotients. The need for *equivalence relations* comes from the fact that many functions can induce the same restriction. The need for *partial* equivalence relations comes from the fact that not every function induces a compatible restriction.

Plotkin [1980] and Statman [1985] have characterized the elements of $\bigcup_\alpha A_\alpha$ that are interpretations of terms under a given environment ρ.

Exercises 10.3.7 a) *If A has smaller cardinality than B, then $\mathcal{F}_B$ distinguishes all the $\beta\eta$-different terms already distinguished by $\mathcal{F}_A$.* (Berardi) (Hint: modify the proof of 10.3.6 to show that if $\mathcal{F}_B$ identifies two terms, then so does $\mathcal{F}_A$.)

b) *Two $\beta\eta$-different terms can be distinguished in $\mathcal{F}_A$, for some finite set A.* (Statman [1980]) (Hint:)

Thus we can restrict attention to the models $\mathcal{F}_{\{0,\dots,n\}}$ and $\mathcal{F}_\omega$, and the latter can be seen as the limit of the family $\{\mathcal{F}_{\{0,\dots,n\}}\}_{n\in\omega}$.

The Completeness Theorem shows that $\beta\eta$-different terms can be distinguished in appropriate functional models, by giving them different interpretations. Of course, we cannot expect to improve the result to hold for β-different terms. On the one hand, the functional models are extensional, and thus cannot distinguish $\beta\eta$-equal terms. On the other hand, there are β-different terms that behave extensionally the same way, i.e. are $\beta\eta$-equal.

10.4 Categorical Models

By the Lawvere-Lambek and Currry-Howard isomorphisms,

proofs of logic = terms of λ-calculus = morphisms of cartesian closed categories.

So c.c.c.'s should be models of the Typed Lambda Calculus.

****

Prove Soundness and Completeness Theorems, as for functional models.

****

Main observation: in the functional models we used only *eval* and *curry*, so the result should extend to any c.c.c.

****

Full functional models provide an interesting class, with regularity properties expressed by the Soundness and Completeness Theorem, but with one drawback: the absurd cardinality of the sets involved. Indeed, at every arrow type we introduce a power set (the set of all functions from a given set to another). If A is countably infinite, then $\mathcal{F}_A$ is actually comparable with the following segment of the cumulative hierarchy of Set Theory:

$$\begin{array}{lcl} V_0 & = & \emptyset \\ V_{n+1} & = & \mathcal{P}(V_n) \\ V_\omega & = & \bigcup_{n\in\omega} V_n \\ V_{\omega+n+1} & = & \mathcal{P}(V_{\omega+n}) \\ V_{\omega+\omega} & = & \bigcup_{n\in\omega} V_{\omega+n}. \end{array}$$

It is certainly no surprise that $V_{\omega+\omega}$ is a model of the Typed Lambda Calculus, since it is already a model of a substantial fragment of Set Theory.[2] In particular, the elements used to interpret the countably many terms are only a fraction of the elements available, and the model is highly redundant.

We try now to cut down such a model to a more manageable one, in which fewer elements are introduced. This is parallel to the procedure in Set Theory, in which thinner hierarchies are obtained by replacing the full power set operation $\mathcal{P}$ by restricted versions of it (e.g. taking not every subset of a given set, but only the ones 'definable' in an appropriate language).

The main idea is well illustrated by considering functions on real numbers. There are of course lots of functions from reals to reals, since each function is determined by its full graph

$$\{(x, f(x)) : x \in \mathcal{R}\}$$

[2] More precisely, the axioms of Zermelo [1908], i.e. the usual ones without replacement, but including infinity.

But there aren't too many *continuous* functions, since each continuous function f is actually determined by the partial graph

$$\{(r, f(r)) : r \in \mathcal{Q}\},$$

i.e. by its values on the rationals: indeed, for every $x \in \mathcal{R}$,

$$f(x) = \lim_{r \in \mathcal{Q} \wedge r < x} f(r)$$

In particular, while the functions are more than the reals, there are only as many countinuous functions as there are reals.[3]

We now try to formulate the notion of continuity in more generality. Of course, a whole branch of mathematics (called *topology*) deals with such a problem. Here we only need a special version of continuity, and at present we have no need for general topological notions. For our purposes, the following notion is sufficiently general and simple.

Categorical models

We can now apply the previous concepts to the construction of models of the Typed Lambda Calculus. The idea is to proceed as in the case of the full functional models, with the following differences: we start not with any set, but with any *c.c.p.o.*; and we proceed by taking not all functions, but only the *continuous* ones.

Definition 10.4.1 *For any object D of a cartesian closed category $\mathcal{C}$, the* **categorical type hierarchy over D** *is the structure $\{D_\alpha\}_\alpha$ defined as follows:*

1. *$D_\alpha = D$, for α atomic*
2. *$D_{\alpha \to \beta} = D_\alpha \Rightarrow D_\beta$.*

The definition of the model becomes then the following.

Definition 10.4.2 *The* **categorical model over D in $\mathcal{C}$** *is the structure*

$$\mathcal{C}_D = \langle \{D_\alpha\}_\alpha, [\![\ ]\!]^{\mathcal{C}_D} \rangle,$$

where $[\![\ ]\!]^{\mathcal{C}_D}$ is the canonical interpretation over $\{D_\alpha\}_\alpha$.

The next result shows that typed λ-terms can be interpreted not only as functions, but as morphisms in any cartesian closed category.

[3] This can be made precise by using the notion of cardinality in Set Theory: there are $2^{\aleph_0}$ reals, and hence $2^{2^{\aleph_0}}$ functions on the reals; but there are only $\aleph_0$ rationals, and hence only $2^{\aleph_0}$ continuous functions on the reals, i.e. as many as the reals themselves.

Theorem 10.4.3 Soundness for Categorical Models. *For any object D in a cartesian closed category $\mathcal{C}$, the structure $\mathcal{C}_D$ is an extensional model of the Typed Lambda Calculus.*

Proof. The proof consists of two parts: to show that the canonical interpretation is well-defined, and to show that $\beta\eta$-equality is preserved. The second part is literally the same as the proof of 10.3.4, and we thus concentrate on the first one.

****

Replace everything in terms of *id*, *eval*, *curry*, and use only the adjointness property.

Continuous = morphism

****

We prove, by induction on t, that $[\![t^\alpha]\!]_\rho \in D_\alpha$ for any environment ρ on $\{D_\alpha\}_\alpha$. Recall that, by definition 10.3.1,

$$[\![t^\alpha]\!]_\rho = \begin{cases} \rho(x^\alpha) & \text{if } t^\alpha = x^\alpha \\ [\![u^{\gamma\to\alpha}]\!]_\rho([\![v^\gamma]\!]_\rho) & \text{if } t^\alpha = u^{\gamma\to\alpha}v^\gamma \\ \Lambda X^\gamma.\, [\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]} & \text{if } t^\alpha = \lambda x^\gamma.\, u^\beta, \end{cases}$$

where $\Lambda X^\gamma.\, [\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]}$ denotes the function

$$a \in D_\gamma \longmapsto [\![u^\beta]\!]_{\rho[x^\gamma:=a]} \in D_\beta.$$

In the first case, $[\![t^\alpha]\!]_\rho \in D_\alpha$ by definition of environment. In the second case, $[\![u^{\gamma\to\alpha}]\!]_\rho$ is an element of $D_{\gamma\to\alpha}$, and hence a function from D_γ to D_α, by the induction hypothesis. Then $[\![t^\alpha]\!]_\rho \in D_\alpha$, by the induction hypothesis $[\![v^\gamma]\!]_\rho \in D_\gamma$. In the last case, $[\![u^\beta]\!]_\rho \in D_\beta$ by the induction hypothesis, and thus $[\![t^\alpha]\!]_\rho$ is a function from D_γ to D_β. The whole point is to show that it is a *continuous* function, since then it is in $D_{\gamma\to\beta} = D_\alpha$.

We prove, by induction on u, that $[\![u^\beta]\!]_{\rho[x^\alpha:=X^\alpha]}$ *is a continuous function of* X^α*, for any type* α. In the course of the proof we will discover that we need a number of facts, which we will prove separately at the end.

If u^β is a variable, there are two cases: if $\beta = \alpha$ and $u^\beta = x^\alpha$ then

$$[\![u^\beta]\!]_{\rho[x^\alpha:=X^\alpha]} = [\![x^\alpha]\!]_{\rho[x^\alpha:=X^\alpha]} = X^\alpha$$

and obviously, for every α, the *identity function* on D_α is continuous; if $u^\beta = y^\beta$ with $y^\beta \neq x^\alpha$ then

$$[\![u^\beta]\!]_{\rho[x^\alpha:=X^\alpha]} = [\![y^\beta]\!]_{\rho[x^\alpha:=X^\alpha]} = \rho(y^\beta)$$

does not depend on X^α and obviously, for every α and β, any *constant function* from D_α to D_β is continuous.

If $u^\beta = u_1^{\gamma\to\beta} u_2^\gamma$ then

$$[\![u^\beta]\!]_{\rho[x^\alpha := X^\alpha]} = [\![u_1^{\gamma\to\beta} u_2^\gamma]\!]_{\rho[x^\alpha := X^\alpha]} = [\![u_1^{\gamma\to\beta}]\!]_{\rho[x^\alpha := X^\alpha]}([\![u_2^\gamma]\!]_{\rho[x^\alpha := X^\alpha]}).$$

By the induction hypothesis, both $[\![u_1^{\gamma\to\beta}]\!]_{\rho[x^\alpha := X^\alpha]}$ and $[\![u_2^\gamma]\!]_{\rho[x^\alpha := X^\alpha]}$ are continuous functions of X^α. For every γ and β, we let the *application function*

$$\mathbf{App}_{\gamma,\beta} : [D_\gamma \to D_\beta] \times D_\gamma \to D_\beta$$

be defined, for $f^{\gamma\to\beta} \in [D_\gamma \to D_\beta]$ and $a \in D_\gamma$, as

$$\mathbf{App}_{\gamma,\beta}(f^{\gamma\to\beta}, a) = f^{\gamma\to\beta}(a).$$

Then, since

$$[\![u^\beta]\!]_{\rho[x^\alpha := X^\alpha]} = \mathbf{App}_{\gamma,\beta}([\![u_1^{\gamma\to\beta}]\!]_{\rho[x^\alpha := X^\alpha]}, [\![u_2^\gamma]\!]_{\rho[x^\alpha := X^\alpha]}),$$

we will need to show that $\mathbf{App}_{\gamma,\beta}$ is continuous, and that the *composition* of continuous functions is still continuous (the latter condition is needed because $[\![u]\!]$ is the composition of $\mathbf{App}$ with $[\![u_1]\!]$ and $[\![u_2]\!]$).

Finally, if $\beta = \gamma \to \delta$ and $u^\beta = \lambda y^\gamma.\, u_1^\delta$, then

$$[\![u^\beta]\!]_{\rho[x^\alpha := X^\alpha]} = [\![\lambda y^\gamma.\, u_1^\delta]\!]_{\rho[x^\alpha := X^\alpha]} = \Lambda Y^\gamma.\, [\![u_1]\!]_{\rho[x^\alpha := X^\alpha; y^\gamma := Y^\gamma]}.$$

By the induction hypothesis, $[\![u_1]\!]_{\rho[x^\alpha := X^\alpha; y^\gamma := Y^\gamma]}$ is a continuous function of X^α and Y^γ, separately. Thus we need to know that it is *continuous in both variables* simultaneously, and that the Λ-*abstraction* w.r.t. one variable of a function continuous in two variables yields a continuous function of the remaining variable (we also need the continuity of composition, but this has already been considered above).

We now prove all these separate facts, in some generality. This will require a few preliminary steps.

****

The next result is the analogue of 10.3.6.

Theorem 10.4.4 Completeness for Continuous Functional Models (Plotkin [1980]) *For any infinite c.c.p.o.* $(D, \sqsubseteq)$,

$$t_1^\alpha =_{\beta\eta} t_2^\alpha \;\Leftrightarrow\; \mathcal{C}_D \models t_1^\alpha = t_2^\alpha.$$

How general is this? for which cartesian closed categories does it hold?

****

Completeness Theorem: every model of typed λ-calculus is a categorical model. Comment on categories as:

- syntactically equivalent to lambda calculus via the isomorphisms (equational theory)
- semantically adequate as models (the examples of c.c.c.)

10.5 Semantical Proof of Strong Normalization ⋆

We have proved the Strong Normalization Theorem in 8.3.2, by means of a syntactical proof which was appropriate in two respects. First, the result itself is syntactical, and hence there is no apparent need to use semantical notions to prove it. Second, the particular proof we used is quite flexible, and is applicable to other type systems.

We are now going to give a semantical proof which, although not as flexible, will provide a clearer picture of what is going on, thus adding an element of intuition to a crucial phenomenon.

The main idea is to use the semantical notions to build a model not of β-equality, but rather of β-reduction.

Well-founded relations and monotone functions

We define a type hierarchy $\{N_\alpha\}_\alpha$ over the set of natural numbers, with *well-founded* partial orderings $<_\alpha$ propagated from the usual ordering of natural numbers, and by taking at arrow levels not *all* functions, but only the *strictly monotone* ones. Well-foundedness and monotonicity are the main ideas of the proof, to be exploited as follows.

We define the interpretation of terms in such a way that the interpretation of a redex is strictly greater than the interpretation of its reduct. By hereditary monotonicity, the interpretation of a term will be strictly greater than the interpretation of any term obtained from it by one application of β-reduction. By well-foundedness, the interpretation of a term can decrease only finitely often. Thus only finitely many successive β-reductions can be performed, starting from a given term. This is exactly what Strong Normalization claims.

We start by defining the type hierarchy over N, together with the associated partial orderings.

Definition 10.5.1 (Gandy [1980]) *The* **Gandy type hierarchy** *is the structure* $\{(N_\alpha, <_\alpha)\}_\alpha$ *defined as follows:*

1. $N_\alpha = N$ *and* $<_\alpha = <$ *for* α *atomic*
2. $N_{\alpha\to\beta}$ *is the set of all strictly monotone functions from* N_α *to* N_β*, i.e.*

$$N_{\alpha\to\beta} = \{f \in (N_\beta)^{N_\alpha} : (\forall a, b \in N_\alpha)(a <_\alpha b \ \Rightarrow \ f(a) <_\beta f(b))\},$$

ordered by:

$$f <_{\alpha\to\beta} g \ \Leftrightarrow \ (\forall a \in N_\alpha)(f(a) <_\beta g(a)).$$

Note that $f <_{\alpha\to\beta} g$ if and only if the graph of f is pointwise below the graph of g, in the ordering $<_\beta$.

We first consider the canonical interpretation over $\{N_\alpha\}_\alpha$ defined in 10.3.1. Given an environment ρ on $\bigcup_\alpha N_\alpha$, i.e. a function assigning to every variable of type α an element of N_α, then

$$[\![t^\alpha]\!]_\rho = \begin{cases} \rho(x^\alpha) & \text{if } t^\alpha = x^\alpha \\ [\![u^{\gamma\to\alpha}]\!]_\rho([\![v^\gamma]\!]_\rho) & \text{if } t^\alpha = u^{\gamma\to\alpha}v^\gamma \\ \Lambda X^\gamma.\, [\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]} & \text{if } t^\alpha = \lambda x^\gamma.\, u^\beta, \end{cases}$$

where $\Lambda X^\gamma.\, [\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]}$ denotes the function

$$a \in N_\gamma \longmapsto [\![u^\beta]\!]_{\rho[x^\gamma:=a]} \in N_\beta.$$

Unfortunately, this definition does *not* provide an interpretation as it stands, since in the last clause x^γ may not occur free in u^β, in which case $\Lambda X^\gamma.\, [\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]}$ is a *constant* function, hence not a member of $N_{\gamma\to\beta}$ (which contains only strictly monotone functions from N_γ to N_β).

We thus have to modify the canonical interpretation, by ensuring that we always obtain strictly monotone functions. The idea is to combine, in a way preserving monotonicity, $[\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]}$ with some expression depending explicitly on X^γ. The simplest such expression is of course X^γ itself, but the problem is that $[\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]}$ describes an element of N_β, while X^γ describes an element of N_γ. This is easily taken care of. It is enough to choose any function $L_{\gamma\to\beta} \in N_{\gamma\to\beta}$, and to project X^γ over N_β, by using $L_{\gamma\to\beta}(X^\gamma)$ in a monotone way. This requires the existence of $L_{\gamma\to\beta}$, which we will have to prove later. For the rest of the argument, the particular form of $L_{\gamma\to\beta}$ is inessential.

The second problem is how to combine pairs of elements of N_β into a single one, in a monotone way. At atomic levels, where the underlying structure is the set N of natural numbers, this is easily obtained by considering the sum function, which is monotone in both arguments. It is then enough to propagate such a function at every level.

1. $+_\alpha = +$ for α atomic

2. $+_{\alpha\to\beta}$ is obtained by adding values pointwise, i.e. by defining

   $$(f +_{\alpha\to\beta} g)(a) \overset{\text{def}}{=} f(a) +_\beta g(a)$$

   for $f, g \in N_{\alpha\to\beta}$ and $a \in N_\alpha$.

Given an environment ρ on $\bigcup_\alpha N_\alpha$, we now define:

$$[\![t^\alpha]\!]_\rho = \begin{cases} \rho(x^\alpha) & \text{if } t^\alpha = x^\alpha \\ [\![u^{\gamma\to\alpha}]\!]_\rho([\![v^\gamma]\!]_\rho) & \text{if } t^\alpha = u^{\gamma\to\alpha}v^\gamma \\ \Lambda X^\gamma.\, ([\![u^\beta]\!]_{\rho[x^\gamma:=X^\gamma]} +_\beta L_{\gamma\to\beta}(X^\gamma)) & \text{if } t^\alpha = \lambda x^\gamma.\, u^\beta. \end{cases}$$

It is now an immediate, by induction, that $[\![t^\alpha]\!]_\rho \in N_\alpha$, because

$$[\![u^\beta]\!]_{\rho[x^\gamma := X^\gamma]} +_\beta L_{\gamma\to\beta}(X^\gamma)$$

is monotone in X^γ.

At this point it can be proved quite easily, by induction on u, that $[\![u]\!]_{\rho[x:=X]}$ is a non-decreasing function of X. Thus the interpretation of a redex will automatically be greater than or equal to the interpretation of its reduct. Indeed,

$$\begin{array}{rcl}
[\![(\lambda x^\alpha.\, u^\beta)v^\alpha]\!]_\rho & = & [\![\lambda x^\alpha.\, u^\beta]\!]_\rho([\![v^\alpha]\!]_\rho) \\
& = & (\Lambda X^\alpha.\, [\![u^\beta]\!]_{\rho[x^\alpha := X^\alpha]} +_\beta L_{\alpha\to\beta}(X^\alpha))([\![v^\alpha]\!]_\rho) \\
& = & [\![u^\beta]\!]_{\rho[x^\alpha := [\![v^\alpha]\!]_\rho]} +_\beta L_{\alpha\to\beta}([\![v^\alpha]\!]_\rho) \\
& = & [\![u^\beta[x^\alpha := v^\alpha]]\!]_\rho +_\beta L_{\alpha\to\beta}([\![v^\alpha]\!]_\rho) \\
& \geq_\beta & [\![u^\beta[x^\alpha := v^\alpha]]\!]_\rho .
\end{array}$$

However, in general we do not have a strict inequality. E.g., if α and β are both atomic, then $[\![v^\alpha]\!]_\rho$ might be 0, and $L_{\alpha\to\beta}$ might be the identity (actually, this is how it is defined below). Then $L_{\alpha\to\beta}([\![v^\alpha]\!]_\rho) = 0$.

We thus need to modify our definition of interpretation one last time, by using in the last clause a strictly increasing function. At atomic levels, where the underlying structure is the set N of natural numbers, this is easily obtained by considering e.g. the successor function S. It is then enough to propagate such a function at every level.

1. $S_\alpha = S$ for α atomic
2. $S_{\alpha\to\beta}$ is obtained by increasing values pointwise, i.e. by defining

$$S_{\alpha\to\beta}(f)(a) \overset{\text{def}}{=} S_\beta(f(a))$$

for $f \in N_{\alpha\to\beta}$ and $a \in N_\alpha$.

We are now ready to define the appropriate structure for our proof.

Definition 10.5.2 (Gandy [1980]) *Given the* **Gandy structure**

$$\mathcal{G} = \{(N_\alpha, <_\alpha, +_\alpha, S_\alpha)\}_\alpha,$$

and an environment ρ on it, we define the **Gandy interpretation** $[\![\ \ ]\!]^{\mathcal{G}}_\rho$ *by induction on terms, as follows:*

$$[\![t^\alpha]\!]^{\mathcal{G}}_\rho = \left\{ \begin{array}{ll}
\rho(x^\alpha) & \textit{if } t^\alpha = x^\alpha \\
[\![u^{\gamma\to\alpha}]\!]^{\mathcal{G}}_\rho([\![v^\gamma]\!]^{\mathcal{G}}_\rho) & \textit{if } t^\alpha = u^{\gamma\to\alpha}v^\gamma \\
\Lambda X^\gamma.\, S_\beta([\![u^\beta]\!]^{\mathcal{G}}_{\rho[x^\gamma := X^\gamma]} +_\beta L_{\gamma\to\beta}(X^\gamma)) & \textit{if } t^\alpha = \lambda x^\gamma.\, u^\beta .
\end{array}\right.$$

Theorem 10.5.3 (Gandy [1980]) *The interpretation $[\![t^\alpha]\!]^{\mathcal{G}}$ of a term t^α strictly decreases in $<_\alpha$ when a β-reduction is performed inside it. Precisely, for any environment ρ:*

$$t_1^\alpha \longrightarrow_{1\beta} t_2^\alpha \;\Rightarrow\; [\![t_1^\alpha]\!]_\rho^{\mathcal{G}} >_\alpha [\![t_2^\alpha]\!]_\rho^{\mathcal{G}}.$$

Proof. By induction on the definition of $\longrightarrow_{1\beta}$. For the base case of a redex we have:

$$\begin{array}{lll} [\![(\lambda x^\alpha.\, u^\beta)v^\alpha]\!]_\rho & = & [\![\lambda x^\alpha.\, u^\beta]\!]_\rho([\![v^\alpha]\!]_\rho) \\ & = & (\Lambda X^\alpha.\, S_\beta([\![u^\beta]\!]_{\rho[x^\alpha := X^\alpha]} +_\beta L_{\alpha\to\beta}(X^\alpha)))([\![v^\alpha]\!]_\rho) \\ & = & S_\beta([\![u^\beta]\!]_{\rho[x^\alpha := [\![v^\alpha]\!]_\rho]} +_\beta L_{\alpha\to\beta}([\![v^\alpha]\!]_\rho)) \\ & = & S_\beta([\![u^\beta[x^\alpha := v^\alpha]]\!]_\rho +_\beta L_{\alpha\to\beta}([\![v^\alpha]\!]_\rho)) \\ & >_\beta & [\![u^\beta[x^\alpha := v^\alpha]]\!]_\rho. \end{array}$$

The remaining cases are easily dealt with by induction as in 10.3.4, using monotonicity. □

Corollary 10.5.4 Strong Normalization Theorem. *For every typed λ-term t^α, there is no infinite sequence of reductions starting from t^α.*

Proof. By the theorem, it is enough to show that the partial orderings $<_\alpha$ are well-founded, i.e. there is no infinite descending sequence of elements at any level. Then there can be no infinite descending sequence of interpretations, and hence of β-reductions.

That $<_\alpha$ is well-founded is easily proved by induction on α. For α atomic, it is so for $<$ on N. For $\alpha \to \beta$, any infinite descending sequence of functions

$$f_0 >_{\alpha\to\beta} f_1 >_{\alpha\to\beta} f_2 >_{\alpha\to\beta} \cdots$$

would produce, by definition of $<_{\alpha\to\beta}$, an infinite descending sequence

$$f_0(a) >_\beta f_1(a) >_\beta f_2(a) >_\beta \cdots,$$

where a is any element of N_α. □

We still have to tie one loose string, namely the *existence of elements L_α in every level N_α*. This step of proving that the hierarchy is not trivial corresponds, in the syntactical proof of Strong Normalization, to the step of proving that $\mathcal{C}$ is not trivial (by showing that every variable is in $\mathcal{C}$).

This is easily done by induction on α, as follows (where, for simplicity, 0 indicates any atomic type):

$$\begin{array}{lll} L_0 & = & 0 \in N \\ L_{0\to 0} & = & \Lambda X^0.\, X^0 \\ L_{(\alpha\to\beta)\to 0} & = & \Lambda Y^{\alpha\to\beta}.\, L_{\beta\to 0}(Y^{\alpha\to\beta}(L_\alpha)) \\ L_{\alpha\to(\beta\to\gamma)} & = & \Lambda X^\alpha.\, [\Lambda Y^\beta.\, (L_{\alpha\to\gamma}(X^\alpha) +_\gamma L_{\beta\to\gamma}(Y^\alpha))]. \end{array}$$

It is easy to check, inductively, that the L_α's are strictly monotone functions. Thus they do belong to N_α.

Regarding the formal difficulty of the proof, we cannot yet claim that it is formalizable in First-Order Arithmetic, because the statement that *there is no infinite sequence of reductions* is not even expressible in a first-order language (we need functions to formalize infinite sequences, and we cannot quantify over functions in a first-order language). But an additional step does show that *for every term t^α there is a number n_{t^α} such that every finite proper sequence of reductions starting from t^α must be of length less than n_{t^α}*, and both this statement and its proof can be expressed in First-Order Arithmetic.

The strengthened result can be obtained as follows. From 10.5.3, by monotonicity of $L_{\alpha\to 0}$, we have

$$t_1^\alpha \longrightarrow_{1\beta} t_2^\alpha \;\Rightarrow\; L_{\alpha\to 0}([\![t_1^\alpha]\!]_\rho) > L_{\alpha\to 0}([\![t_2^\alpha]\!]_\rho).$$

It is thus enough to let

$$n_{t^\alpha} = L_{\alpha\to 0}([\![t^\alpha]\!]_\rho)$$

for any fixed assignment ρ, e.g. for

$$\rho(x^\alpha) = L_\alpha.$$

Of course, in a formal proof we need to show that $L_{\alpha\to 0}([\![t^\alpha]\!]_\rho)$ is indeed a number, and this can be proved as in the Weak Normalization Theorem (whose proof is easily formalizable). Thus a formal proof actually consists of two separate steps:

1. a proof that Weak Normalization implies Strong Normalization (10.5.4 plus the additional remarks above)

2. a proof of Weak Normalization (as in 8.3.1).

æ

Chapter 11

Computability

We have introduced the Typed Lambda Calculus as a theory of functions and proved a number of results about it, but we have not said yet how we can actually perform practical computations, for example of numerical functions. The goal of this chapter is to introduce numbers in our theory, and to characterize both its computational power and its computational complexity.

Since we will discover that, in a precise sense, the 'power' of the Typed Lambda Calculus is quite limited, there are good reasons to devise stronger systems. Actually, to reach a computational power equivalent to the usual universal systems, such as the Turing machines, we will have to drop type restrictions completely and step to the Untyped Lambda Calculus.

11.1 Numerals

The idea we exploit to introduce natural numbers in the Typed Lambda Calculus is quite simple. If we think of the *use* of natural numbers, as opposed to their abstract *meaning*, we notice that they appear in everyday life as numbers of something. For example, a request of 'giving three apples' can be interpreted as a request of iterating the operation of 'giving an apple' three times. In the Typed Lambda Calculus, where all we can do is to perform operations, we can thus think of numbers as *exponents* of such operations, telling how many times we have to iterate them.

The notion of **iteration** of a unary function f on a given argument x is defined inductively as:

$$\begin{array}{rcl} f^{(0)}(x) & = & x \\ f^{(n+1)}(x) & = & f(f^{(n)}(x)). \end{array}$$

The first clause is simply a convention, justified by the desire of having the second clause work for all n, including 0. In this case we obtain a single iteration, which correspond to usual applications:

$$f^{(1)}(x) = f(f^{(0)}(x)) = f(x).$$

We can now think of the number n as being an operator that, given a function f and an argument x, produces the n-th iteration $f^{(n)}(x)$ of f on x. Since in the Typed Lambda Calculus all we really have is unary functions, this is a perfectly general intuition, and can be taken as a *definition* of numbers in the theory.

Definition 11.1.1 (Peano [1891], Wittgenstein [1921], Church [1933]) *For any type α, the numeral $\overline{n}$ of type*

$$N_\alpha = (\alpha \to \alpha) \to (\alpha \to \alpha)$$

is the λ-term that produces n iterations of an object of type $\alpha \to \alpha$ on an object of type α, i.e.

$$\overline{n}^{N_\alpha} = \lambda f^{\alpha \to \alpha} x^\alpha . f^{(n)} x,$$

where

$$f^{(n)}x = \underbrace{f(\cdots(f}_{n \text{ times}}\, x)\cdots).$$

To increase readability types can be omitted, when no confusion arises.

Notice that f has to have type $\alpha \to \beta$, if we want to apply it to an argument x of type α. Then fx has type β. If we want to be able to iterate f, i.e. to apply it to fx, then we must have $\beta = \alpha$. This justifies the type $\alpha \to \alpha$ for f.

The way it stands, the definition of numerals is parametrized on the given type α. In particular, there are numerals at every type level N_α. The results that follow will not depend on the particular α we choose to start with, in the sense that they will be uniform in α.

Since numerals are terms in normal form, by the Church-Rosser Theorem they are β-different if they are *syntactically* different. Thus

$$n \neq m \;\Rightarrow\; \overline{n}^{N_\alpha} \neq_\beta \overline{m}^{N_\alpha}.$$

11.2 Representable Functions

Having at hand a version of numbers inside the Typed Lambda Calculus, we can represent a numerical function f by a term F that behaves on the numerals $\overline{n}$ the way f behaves on the numbers n.

Definition 11.2.1 *A n-ary function f is* **representable** *in the Typed Lambda Calculus at level N_α if there is a closed typed term F such that, for every $x_1, \dots,$ x_n and y,*

$$f(x_1, \dots, x_n) = y \;\Leftrightarrow\; F^{N_\alpha \to (\cdots (N_\alpha \to N_\alpha)\cdots)} \overline{x}_1^{N_\alpha} \cdots \overline{x}_n^{N_\alpha} =_\beta \overline{y}^{N_\alpha}.$$

It is important to note that, while α can be any type in the previous definition, it has to be the same for arguments and values. A different definition, usually referred to as **skew representability**, allows different types for arguments and values, i.e.

$$f(x_1, \dots, x_n) = y \;\Leftrightarrow\; F^{N_{\alpha_1} \to (\cdots (N_{\alpha_n} \to N_\beta)\cdots)} \overline{x}_n^{N_{\alpha_1}} \cdots \overline{x}_n^{N_{\alpha_n}} =_\beta \overline{y}^{N_\beta}.$$

The reason we adopt the stricter definition is that it gives better results. In particular, a complete characterization of the skew representable functions in the Typed Lambda Calculus is not known, and does not appear to be any natural class of functions, whereas this is the case for the representable functions.

Examples

We start by representing the basic functions of arithmetic.

Proposition 11.2.2 *The constant functions, as well as sum and product, are representable in the Typed Lambda Calculus.*

Proof. First, the constant function with value n is representable in the natural way, as

$$\lambda x^{N_\alpha}.\, \overline{n}^{N_\alpha}.$$

Second, iterating a function $m+n$ times is the same as iterating it first m times, and then n times. Since

$$f^{(m+n)}(x) = f^{(m)}(f^{(n)}(x)) \quad\Rightarrow\quad \overline{m+n} f x =_\beta \overline{m} f (\overline{n} f x),$$

sum is representable by

$$\lambda m^{N_\alpha} n^{N_\alpha} f^{\alpha \to \alpha} x^\alpha.\, m f (n f x).$$

Notice that, although we are interested in the behavior of this term only on numerals $\overline{m}$ and $\overline{n}$, we do not have variables ranging over numerals, and we thus use variables m and n (not to be confused with numbers!) ranging over all terms. Thus the term just defined can actually be applied to any pair of terms of type N_α, not only to numerals. However, the only *closed* terms of type N_α are the numerals.

Third, iterating a function $m \cdot n$ times is the same as iterating n times its m-fold iteration. Since

$$f^{(m \cdot n)}(x) = (f^{(m)})^{(n)}(x) \quad \Rightarrow \quad \overline{m \cdot n} f x =_\beta \overline{n}(\overline{m} f)x,$$

product is representable by

$$\lambda m^{N_\alpha} n^{N_\alpha} f^{\alpha \to \alpha} x^\alpha .\, n(mf)x. \quad \Box$$

Corollary 11.2.3 *All polynomials with natural numbers as coefficients are representable in the Typed Lambda Calculus.*

Proof. It is enough to notice that the class of representable functions is obviously closed under composition. Indeed, the term representing a composition of functions is the application of the terms representing the composed functions. $\Box$

We could think of continuing in the same vein, and representing the exponential function in a similar way. Indeed, iterating a function m^n times is the same as iterating n times the function that iterates its argument m times. Thus

$$f^{(m^n)}(x) = f^{\overbrace{m \cdots m}^{n \text{ times}}}(x) \quad \Rightarrow \quad \overline{m^n} f x =_\beta (\overline{n}\, \overline{m}) f x.$$

However, this does not produce a permissible representation. If the types of $\overline{m}$, f and x are N_α, $\alpha \to \alpha$ and α, then $\overline{n}$ must have a type $N_\alpha \to N_\alpha = N_{\alpha \to \alpha}$. This goes against the definition of representability, which requires the same type for arguments and values. The nonpermissible term

$$\lambda n^{N_{\alpha \to \alpha}} m^{N_\alpha} f^{\alpha \to \alpha} x^\alpha .\, (nm) f x$$

does show, in the terminology introduced above, that *the exponential function is skew representable*. Obviously, this discussion only shows that the particular representation we had in mind does not work. A more sophisticated argument, given below, is needed to prove that the exponential function is not representable in our sense, i.e. that *no* representation for it exists.

The next result shows that the polynomials do not exhaust the class of representable functions.

Proposition 11.2.4 *Test on zero is representable in the Typed Lambda Calculus.*

Proof. Formally, test on zero is the function δ defined as follows:

$$\delta(x_1, x_2, y) = \begin{cases} x_1 & \text{if } y = 0 \\ x_2 & \text{otherwise} \end{cases}$$

The idea for the representation of δ comes from the following observations.

- Given the numeral $\overline{0}$, then its application to two arguments of the appropriate types gives the second one. Thus, if we want the output to be x_1, we just have to put it as the second argument.
- Given the numeral $\overline{n+1}$, then its application to two arguments of the appropriate types applies the first one at least once. Thus, if we want the output to be x_2, we just have to put the constant function with value x_2 as the first argument.

We are thus led to consider the expression

$$\overline{y}(\lambda z.\, x_2)x_1.$$

The problem is that this is not a term, because the types do not match. Indeed, $\overline{y}$ has a fixed type N_α, and its arguments must then have type $\alpha \to \alpha$ and α, respectively. But if x_1 and x_2 are numerals, they have type N_α. Thus, even if z has type α, the first argument $\lambda z.\, x_2$ of $\overline{y}$ has type $\alpha \to N_\alpha$, and the second argument x_1 has type N_α.

This is easy to fix, by first applying x_1 and x_2 to the arguments $f^{\alpha\to\alpha}$ and x^α, and then abstracting f and x at the end. In other words, since

$$\lambda z^\alpha.\, x_2^{N_\alpha} f^{\alpha\to\alpha} x^\alpha \qquad \text{and} \qquad x_1^{N_\alpha} f^{\alpha\to\alpha} x^\alpha$$

have type $\alpha \to \alpha$ and α, respectively, we can use

$$\lambda x_1^{N_\alpha} x_2^{N_\alpha} y^{N_\alpha} f^{\alpha\to\alpha} x^\alpha.\, y(\lambda z^\alpha.\, x_2 f x)(x_1 f x)$$

for the representation of δ. $\Box$

Piecewise polynomials

We now combine the previous examples into a general result.

Definition 11.2.5 *An n-ary* **piecewise polynomial** *is a function f of the following form:*

$$f(x_1, \ldots, x_n) = p_A \text{ if } (\forall i)_{1 \le i \le n}(x_i = 0 \Leftrightarrow i \in A),$$

where $A \subseteq \{1, \ldots, n\}$, and p_A is a polynomial in the variables x_i such that $i \notin A$.

Thus a piecewise polynomial of n variables is a function defined by 2^n cases, each one of them determined by which variables are 0, and giving rise to a polynomial in the remaining variables. For example, in the case of two variables we have:

$$f(x_1, x_2) = \begin{cases} c & \text{if } x_1 = 0 \wedge x_2 = 0 \\ p_1(x_1) & \text{if } x_1 \neq 0 \wedge x_2 = 0 \\ p_2(x_2) & \text{if } x_1 = 0 \wedge x_2 \neq 0 \\ p_3(x_1, x_2) & \text{if } x_1 \neq 0 \wedge x_2 \neq 0, \end{cases}$$

with c is a constant, and p_1, p_2 and p_3 are polynomials of the indicated variables.

Proposition 11.2.6 *All piecewise polynomials with natural numbers as coefficients are representable in the Typed Lambda Calculus.*

Proof. The work done above is already sufficient to prove the result, once we notice that the piecewise polynomials can actually be defined as the smallest class of functions containing the constant functions, sum, product and δ, and closed under composition.

More precisely, by adding among the initial functions the projection functions

$$\mathcal{I}_i^n(x_1, \ldots, x_n) = x_i$$

for $1 \le i \le n$, then we only need composition in the form

$$f(x_1, \ldots, x_n) = g(h_1(x_1, \ldots, x_n), \ldots, h_m(x_1, \ldots, x_n)),$$

since the projection functions allow trivial manipulations such as interchange, identification and introduction of variables.

It is obvious that $\mathcal{I}_i^n$ is representable by

$$\lambda x_1 \cdots x_n .\, x_i .$$

Moreover, if g, h_1, $\ldots$, h_m are representable by G, H_1, $\ldots$, H_m, respectively, then f is representable by

$$\lambda x_1 \cdots x_n .\, G(H_1 x_1 \cdots x_n) \cdots (H_m x_1 \cdots x_n).$$

We now check that the piecewise polynomials are indeed the functions belonging to the class defined above. In one direction, we notice that all polynomials are easily obtainable by composition from constants, sums and products, possibly by manipulating variables by use of the projections. Moreover, n-ary piecewise polynomials can be defined by a nesting of 2^n case definitions of the form

$$g(x_1, \ldots, x_n, y) = \begin{cases} h_1(x_1, \ldots, x_n) & \text{if } y = 0 \\ h_2(x_1, \ldots, x_n) & \text{otherwise,} \end{cases}$$

each of which is reducible to composition and δ, as follows:

$$g(x_1, \ldots, x_n, y) = \delta(h_1(x_1, \ldots, x_n), h_2(x_1, \ldots, x_n), y).$$

Conversely, every function belonging to the class defined above can be shown to be a piecewise polynomial, by induction on the definition of the class. □

11.3 Nonrepresentable Functions ⋆

Having studied the representable functions, we now turn to the nonrepresentable ones, with the goal of obtaining a complete classification.

Examples

We use a gentle, semantical method to show that certain functions are not representable.

Proposition 11.3.1 (Statman [1982]) *The characteristic function of equality is not representable in the Typed Lambda Calculus.*

Proof. If the characteristic function of equality were representable in the Typed Lambda Calculus, there would be a term E such that, for any pair of natural numbers n and m,

$$E\,\overline{n}^{N_\alpha}\overline{m}^{N_\alpha} =_\beta \begin{cases} \overline{1}^{N_\alpha} & \text{if } n = m \\ \overline{0}^{N_\alpha} & \text{otherwise,} \end{cases}$$

where $\overline{0}^{N_\alpha} = \lambda f^{\alpha\to\alpha}x^\alpha.\,x$ and $\overline{1}^{N_\alpha} = \lambda f^{\alpha\to\alpha}x^\alpha.\,yx$. For simplicity of notation, we drop the types in the following.

Given any model of the Typed Lambda Calculus, from the representability of E it would follow that

$$n \neq m \;\Rightarrow\; E\,\overline{n}\,\overline{m} =_\beta \overline{0} \;\Rightarrow\; [\![E]\!]([\![\overline{n}]\!], [\![\overline{m}]\!]) = [\![\overline{0}]\!],$$

and

$$n = m \;\Rightarrow\; E\,\overline{n}\,\overline{m} =_\beta \overline{1} \;\Rightarrow\; [\![E]\!]([\![\overline{n}]\!], [\![\overline{m}]\!]) = [\![\overline{1}]\!].$$

Consider now the full functional model over a two-element set $A = \{a, b\}$ defined in 10.3.3, and let α be an atomic type. Since $N_\alpha = (\alpha \to \alpha) \to (\alpha \to \alpha)$ and A has two elements, there are only $(2^2)^{(2^2)} = 256$ elements of type N_α. Since there are infinitely many numerals, there must exist distinct natural numbers n and m with coinciding interpretations $[\![\overline{n}]\!]$ and $[\![\overline{m}]\!]$. It follows from the above that

$$[\![\overline{0}]\!] = [\![E]\!]([\![\overline{n}]\!], [\![\overline{m}]\!]) = [\![E]\!]([\![\overline{n}]\!], [\![\overline{n}]\!]) = [\![\overline{1}]\!],$$

because $n \neq m$, $[\![\overline{n}]\!] = [\![\overline{m}]\!]$ and $n \neq n$. Thus $[\![\overline{0}]\!] = [\![\overline{1}]\!]$.

But the model uses the canonical interpretation of terms, i.e.

$$[\![\overline{0}]\!] = \Lambda F \Lambda X.\,X \qquad \text{and} \qquad [\![\overline{1}]\!] = \Lambda F \Lambda X.\,F(X).$$

It is then enough to choose a function F such that $F(a) = b$, to have

$$[\![\overline{0}]\!](F, a) = a \neq b = [\![\overline{1}]\!](F, a),$$

i.e. $[\![\overline{0}]\!] \neq [\![\overline{1}]\!]$. □

The proof shows that in the Typed Lambda Calculus equality is not representable, because there are nontrivial finite models. The same proof shows that in the Untyped Lambda Calculus there are no nontrivial finite models, because equality *is* representable (see 13.2.3).

Corollary 11.3.2 *The characteristic function of the ordering, as well as predecessor and subtraction, are not representable in the Typed Lambda Calculus.*

Proof. Suppose $\leq$ were representable. Then so would be $=$, because

$$x = y \ \Leftrightarrow \ x \leq y \ \wedge \ y \leq x,$$

and $\wedge$ is representable because so is product (11.2.2).

Suppose subtraction were representable. Then so would be $\leq$, because

$$x \leq y \ \Leftrightarrow \ x - y = 0 \ \Leftrightarrow \ \delta(x-y) = 1,$$

and δ is representable (11.2.4).

Suppose predecessor were representable by P. Then so would be subtraction, because

$$\overline{m-n} = \overline{n} P \, \overline{m},$$

i.e. $m - n$ is the n-th iteration of the predecessor on m. □

Exercises 11.3.3 a) *The model of the Lambda Calculus used in the proof of 11.3.1 has exactly three distinct interpretations of numerals. Actually, $[\![\overline{n}]\!] = [\![\overline{1}]\!]$ for any odd n, and $[\![\overline{n}]\!] = [\![\overline{2}]\!]$ for any even $n \neq 0$.* (Hint: there are only four possible functions f on $\{a, b\}$, and all of them satisfy $f^{(3)} = f$. To prove that $[\![\overline{0}]\!] \neq [\![\overline{2}]\!]$, use the constant function with value b. To prove that $[\![\overline{1}]\!] \neq [\![\overline{2}]\!]$, use the function that interchanges a and b.)

b) *There is a model of the Typed Lambda Calculus with exactly two distinct interpretations of numerals. Actually, $[\![\overline{n}]\!] = [\![\overline{1}]\!]$ for any $n \geq 1$.* (Hint: consider the model obtained by starting with $\{a, b\}$ with the order $a \sqsubset b$ at atomic types, and taking only monotone or constant functions at arrow types. Then there are only three possible such functions f on $\{a, b\}$, and all of them satisfy $f^{(2)} = f$.)

c) *In any nontrivial model of the Typed Lambda Calculus, $[\![\overline{n}]\!] \neq [\![\overline{0}]\!]$ for any $n \geq 1$.* (Hint: otherwise, test on zero would not be representable.)

d) *The theory of $\beta\eta$-equality is not maximal.* (Hint: the equation $\overline{1} = \overline{2}$ is not derivable because $\overline{1} \neq_{\beta\eta} \overline{2}$, but it is consistent by part b).)

e) *There is no maximal extension of the theory of $\beta\eta$-equality.* (Hint: the equation $\overline{0} = \overline{2}$ is not derivable, but it is consistent. By part d), any maximal theory would make $\overline{0} = \overline{1}$, contradicting part c).)

The Characterization Theorem

We now prove the converse of 11.2.6, thus characterizing the class of functions representable in the Typed Lambda Calculus. Unlike in the previous examples, we are here forced to use a brutal, syntactical method.

Theorem 11.3.4 (Schwichtenberg [1975], Statman) *If a function is representable in the Typed Lambda Calculus, then it is a piecewise polynomial.*

Proof. We prove the result only for the case of unary functions, the case of n-ary functions being more cumbersome but similar. Let f be a unary function and F be a closed typed term such that, for some α and every n and m,

$$f(n) = m \Leftrightarrow F^{N_\alpha \to N_\alpha} \overline{n}^{N_\alpha} =_\beta \overline{m}^{N_\alpha}.$$

By the Normalization Theorem, we may suppose that F is in normal form. Since it has type $N_\alpha \to N_\alpha$, it must be either of the form

$$\lambda x^{N_\alpha}.\, x^{N_\alpha},$$

in which case f is the identity function, or of the form

$$F = \lambda x^{N_\alpha} f^{\alpha\to\alpha}.\, u^{\alpha\to\alpha}$$

for some u, in which case

$$F\overline{n} =_\beta \lambda f^{\alpha\to\alpha}.\, u[x := \overline{n}].$$

If $n = 0$, then the effect of substituting $\overline{n}^{N_\alpha}$ for x^{N_α} inside u is to replace every subterm of the form $x^{N_\alpha} v^{\alpha\to\alpha}$ by the term $\overline{0}v$. Since $\overline{0} = \lambda g^{\alpha\to\alpha} z^\alpha . z^\alpha$,

$$\overline{0} v^{\alpha\to\alpha} =_\beta \lambda z^\alpha .\, z^\alpha = \mathbf{I}^{\alpha\to\alpha},$$

which does not depend on n. Thus $u[x := \overline{n}]$ is a constant term. By the hypothesis on F, the normal form of $F\overline{n}$ is a numeral. Then it is equal to $\overline{c}$, for some c.

If $n \neq 0$, we analyze the structure of u, which is a term in normal form of type $\alpha \to \alpha$, with $f^{\alpha\to\alpha}$ and x^{N_α} as the only possible free variables, and α, $\alpha \to \alpha$ and N^α as the only possible types of subterms. We will discover that in this case the normal form of $F\overline{n}$ is a numeral $\overline{p(n)}$ for some polynomial p, so that in the end

$$F\overline{n} =_\beta \begin{cases} \overline{c} & \text{if } \overline{n} =_\beta \overline{0} \\ \overline{p(n)} & \text{if } \overline{n} \neq_\beta \overline{0}. \end{cases}$$

Then

$$f(n) = \begin{cases} c & \text{if } n = 0 \\ p(n) & \text{if } n \neq 0, \end{cases}$$

and f is a piecewise polynomial. The analysis proceeds in two steps:

1. *inductive characterization of all terms in normal form of type $\alpha \to \alpha$, with $f^{\alpha\to\alpha}$ and x^{N_α} as the only possible free variables, and α, $\alpha \to \alpha$ and N_α as the only possible types of subterms*

   We first note that the class $\mathcal{N}$ of all *terms in normal form* can be inductively generated as follows,with all terms of the appropriate types:

- all variables are in $\mathcal{N}$
- if z is a variable and $v_1, \ldots, v_n \in \mathcal{N}$, then $zv_1 \cdots v_n \in \mathcal{N}$
- if $v \in \mathcal{N}$, then $\lambda y.\, v \in \mathcal{N}$.

Indeed, this is simply a restatement of the inductive definition of the class of all terms, with the additional constraint that in an application of two terms, the first one cannot start with a λ. Then, inductively, it must be either a variable or an application that does not start with a λ. In finitely many steps we reduce to variables, and hence to the second clause.

We consider the class $\mathcal{U}$ of all *terms in normal form of type $\alpha \to \alpha$, with $f^{\alpha\to\alpha}$, x^{N_α} and any z^α as the only possible free variables, and α, $\alpha \to \alpha$ and N_α as the only possible types of subterms*. Notice that u belongs to $\mathcal{U}$.[1] We now prove that $\mathcal{U}$ can be inductively generated as follows:

(a) for any y^α and z^α, $\lambda y^\alpha.\, z^\alpha \in \mathcal{U}$

(b) $f^{\alpha\to\alpha} \in \mathcal{U}$

(c) if $v^{\alpha\to\alpha} \in \mathcal{U}$, then $x^{N_\alpha} v^{\alpha\to\alpha} \in \mathcal{U}$

(d) if $v_1, \ldots, v_m \in \mathcal{U}$ are not λ-abstractions, then $\lambda y^\alpha.\, v_1 \cdots (v_m z^\alpha) \in \mathcal{U}$.

Every term inductively generated by the previous clauses is obviously in $\mathcal{U}$. To prove the converse, we consider any term $t^{\alpha\to\alpha} \in \mathcal{U}$ and proceed by induction on the previous characterization of $\mathcal{N}$.

- If $t^{\alpha\to\alpha}$ is a variable, it must be $f^{\alpha\to\alpha}$, x^{N_α} or some z^α. Since only $f^{\alpha\to\alpha}$ has the right type, we are in case (b).
- If $t^{\alpha\to\alpha}$ is of the form $zv_1 \cdots v_n$, with $v_1, \ldots, v_n$ in normal form, then z must either be one of $f^{\alpha\to\alpha}$ and x^{N_α}, or have type α.
  - In the first case, $t = f$. Then we are in case (b).
  - In the second case, $t = x^{N_\alpha} v^{\alpha\to\alpha}$. Since $f^{\alpha\to\alpha}$, x^{N_α} and any z^α are the only possible free variables in t, and α, $\alpha \to \alpha$ and N_α are the only possible types of subterms of t, the same is true for v. Then $v \in \mathcal{U}$, and we are in case (c).
  - The third case cannot occur. Indeed, z^α cannot be applied to any term, because every subterm of a term in $\mathcal{U}$ must have type α, $\alpha \to \alpha$ or N_α. Then z^α should occur alone. But it cannot, because it has the wrong type.

[1] u actually belongs to the class of terms of type $\alpha \to \alpha$ in normal form, with $f^{\alpha\to\alpha}$ and x^{N_α} as the only possible free variables, and α, $\alpha \to \alpha$ and N_α as the only possible types of subterms. The definition of $\mathcal{U}$, which also admits free variables of type α, provides a stronger induction hypothesis needed to make the inductive characterization below possible (see Note 2).

- If $t^{\alpha\to\alpha}$ is of the form $\lambda y^\alpha.\, v^\alpha$, with v in normal form, then v is a term of type α in which only x^{N_α}, $f^{\alpha\to\alpha}$ and variables of type α can occur free. By induction on the characterization of $\mathcal{N}$, we have the following possibilities.

  If v^α is a variable, it must be some z^α, because $f^{\alpha\to\alpha}$ and x^{N_α} have the wrong type. Then we are in case (a).

  If v^α is of the form $zv_1\cdots v_n$, with all v_i in normal form, then every v_i must have type α, $\alpha\to\alpha$ or N_α. Moreover, z must either be one of $f^{\alpha\to\alpha}$ and x^{N_α}, or have type α.

  - In the first case, $v^\alpha = f^{\alpha\to\alpha}v_1^\alpha$. But $f^{\alpha\to\alpha}\in\mathcal{U}$, and v_1^α is a term of type α in which only $f^{\alpha\to\alpha}$, x^{N_α} and variables of type α can occur free. Then we can start again on it.
  - In the second case, $v^\alpha = x^{N_\alpha}v_1^{\alpha\to\alpha}v_2^\alpha$, and only $f^{\alpha\to\alpha}$, x^{N_α} and any z^α can occur free in v_1 and v_2. Thus v_1, which has type $\alpha\to\alpha$, is in $\mathcal{U}$.[2] Hence $x^{N_\alpha}v_1^{\alpha\to\alpha}$ is in $\mathcal{U}$ as well. Since v_2 has type α, we can start again on it.
  - In the third case, it must be $n=0$. Then $v^\alpha = z^\alpha$ and $t = \lambda y^\alpha.\, z^\alpha$.

  In the first two cases, we reduce to (d) in finitely many steps. In the last, we are in case (a).

  If v^α is of the form $\lambda y.\, v_1$, then v_1 has type α, $\alpha\to\alpha$ or N_α, and v cannot have type α. So this case cannot occur.

2. *proof by induction on the previous characterization*
   We now prove that if $u\in\mathcal{U}$ and $n\geq 1$, then the normal form of

   $$u[x := \overline{n}]$$

   can only have the following three forms:

   $$f^{\alpha\to\alpha} \qquad \lambda y^\alpha.\, f^{(p(n))}y^\alpha \qquad \lambda y^\alpha.\, f^{(p(n))}z^\alpha,$$

   for some polynomial p of n. Then the normal form of

   $$F\overline{n} = \lambda f^{\alpha\to\alpha}.\, u[x := \overline{n}]$$

   can only have the following three forms:

   $$\lambda f^{\alpha\to\alpha}.\, f^{\alpha\to\alpha} \qquad \lambda f^{\alpha\to\alpha}y^\alpha.\, f^{(p(n))}y^\alpha \qquad \lambda f^{\alpha\to\alpha}y^\alpha.\, f^{(p(n))}z^\alpha.$$

   Since they must all be numerals, the first and third cases (for $z^\alpha \neq y^\alpha$) cannot occur. In the second case, we get $\overline{p(n)}$.

[2] It is here that the extended definition of $\mathcal{U}$ is needed. Without allowing variables of type α to occur free, we could not claim that $v_1\in\mathcal{U}$.

We proceed by induction on the previous characterization of $\mathcal{U}$.

If $u = \lambda y^\alpha .\, z^\alpha$, possibly with $z^\alpha = y^\alpha$, then $u[x := \overline{n}] = \lambda y^\alpha .\, z^\alpha$. This is of the required form, with $p(n) = 0$.

If $u = f^{\alpha \to \alpha}$, then $u[x := \overline{n}] = f^{\alpha \to \alpha}$. This is of the required form.

If $u = x^{N_\alpha} v^{\alpha \to \alpha}$, then

$$u[x := \overline{n}] = \overline{n}(v[x := \overline{n}]).$$

By the induction hypothesis for $v[x := \overline{n}]$, we have three possibilities:

- $v[x := \overline{n}] = f^{\alpha \to \alpha}$. Then $\overline{n} f y = f^{(n)} y$, and so
$$\overline{n} f = \lambda y .\, f^{(n)} y,$$
which is of the required form, with $p(n) = n$.
- $v[x := \overline{n}] = \lambda y^\alpha .\, f^{(p(n))} y^\alpha$. Then $\overline{n}(\lambda y .\, f^{(p(n))} y) y$ produces, by definition, the n-th iteration of the first argument on the second, i.e. $f^{(p(n) \cdot n)} y$, and so
$$\overline{n}(\lambda y .\, f^{(p(n))} y) = \lambda y .\, f^{(p(n) \cdot n)} y$$
- $v[x := \overline{n}] = \lambda y^\alpha .\, f^{(p(n))} z^\alpha$, with $z^\alpha \neq y^\alpha$. Then $\overline{n}(\lambda y .\, f^{(p(n))} z) y$ produces, as above, the n-th iteration of the first argument on the second. But the first argument is a constant function, which is applied at least once because $n \geq 1$. Thus we get $f^{(p(n))} z$, and so
$$\overline{n}(\lambda y .\, f^{(p(n)} z) = \lambda y .\, f^{(p(n))} z.$$

If $u = \lambda y^\alpha .\, v_1 \cdots (v_m z^\alpha)$, then

$$u[x := \overline{n}] = \lambda y^\alpha .\, (v_1[x := \overline{n}]) \cdots ((v_m[x := \overline{n}]) z^\alpha).$$

By the induction hypothesis for the $v_i[x := \overline{n}]$, we have the following possibilities:

- No $v_i[x := \overline{n}]$ is a constant function. Then there are polynomials p_i such that $v_i[x := \overline{n}] = \lambda y .\, f^{(p_i(n))} y$, and
$$u[x := \overline{n}] = \lambda y .\, f^{(p_1(n) + \cdots + p_m(n))} z,$$
where $p_i(n) = 1$ if $v_i[x := \overline{n}] = f$.
- Some $v_i[x := \overline{n}]$ is a constant function $\lambda y .\, f^{(p_i(n))} z$. Let i be the first index such that this happens. Then
$$u[x := \overline{n}] = \lambda y .\, f^{(p_1(n) + \cdots + p_i(n))} z.$$

> The various cases above introduce constants, identities, sums and products in the exponents, thus producing polynomials by induction.

It should be clear from the proof just given that the case of n-ary functions does not present essential differences with the one just dealt with. We only have to consider, for any variable, the various cases in which it is equal to 0 or not, and they give rise to 2^n possibilities. □

Corollary 11.3.5 Characterization of the Representable Functions. *The functions representable in the Typed Lambda Calculus are exactly the piecewise polynomials.*

Proof. By 11.2.6 and 11.3.4. □

11.4 Complexity

The results of the previous sections characterize the computational *power* of the Typed Lambda Calculus, in terms of its representable functions. We now turn to a complementary characterization of the computational *complexity*, in terms of the number of steps needed to perform normalizations. By the Curry-Howard isomorphism, the results transfer automatically from term normalization in the Typed Lambda Calculus to proof normalization in Natural Deduction.

We refer to Odifreddi [1999] for backgroung on Computational Complexity. For our purposes here, it suffices to say that we measure the complexity of a normalization procedure by evaluating the number of steps it takes, in terms of the following fast growing functions:

$$e_0(x) = x \qquad e_{n+1}(x) = x^{e_n(x)} \qquad s(x) = e_x(x).$$

Notice that e_1 is the usual *exponential* function, e_n is a *generalized exponential* function with a fixed stack of n exponents, and s is a *superexponential* function with a variable stack of exponents, depending on the argument.

A lower bound

Our first goal is to determine a lower bound to the number of steps needed to obtain the normal form of a term. The idea is to find, for infinitely many n, a term of length n so complicated that roughly $s(n)$ reductions are needed to reduce it to a normal form. This proves that the complexity of normalization is greater than any generalized exponential, and is at least superexponential.

Proposition 11.4.1 (Statman [1979]) *For infinitely many typed λ-terms, any normalization procedure takes at least a superexponential number of steps.*

Proof. By the definition of numerals as iterators (11.1.1), we have

$$\overline{n} f x =_\beta f^{(n)} x = f^{(e_0(n))} x.$$

From the example of skew representability on p. 232, we have

$$(\overline{n}\,\overline{n}) f x =_\beta f^{(n^n)} x = f^{(e_1(n))} x.$$

By induction, we similarly obtain

$$\underbrace{(\overline{n} \cdots \overline{n})}_{m+1 \text{ times}} \; f x =_\beta f^{(e_m(n))} x$$

and

$$\underbrace{(\overline{n} \cdots \overline{n})}_{n+1 \text{ times}} \; f x =_\beta f^{(e_n(n))} x = f^{(s(n))} x.$$

Obviously, the types of the various $\overline{n}$ are all distinct.

We now compute how many steps are needed to perform the normalization

$$\underbrace{(\overline{n} \cdots \overline{n})}_{n+1 \text{ times}} \; f x \; \longrightarrow_\beta \; f^{(s(n))} x = \underbrace{f \cdots f}_{s(n) \text{ times}} \; x.$$

Notice that, since $\overline{n}$ has approximately length n, the left-hand-side has length of order n^2. The right-hand-side has instead length of order $s(n)$.

The main observation is that if a term t has length l, then any subterm of t has length at most l, and it has at most l occurrences of any given variable. Since any application of the β-rule to a subterm $(\lambda x.\, u)v$ of t replaces all occurrences of the variable x in the subterm u by the subterm v, it can at most square the length of the given term t.

We must thus go from a term having length of order n^2 to a term having length of order $s(n)$, by a series of steps that can at most square the length:

$$n^2 \qquad (n^2)^2 = n^{2^2} \qquad (n^{2^2})^2 = n^{2^3} \qquad \cdots \qquad n^{2^m} \qquad \cdots$$

To get a value n^{2^m} of the order of $s(n) = e_n(n)$, i.e. with a stack of n exponents, we need an m of the order of $e_{n-2}(n)$, i.e. with a stack of $n-2$ exponents. In particular, we need at least $s(n-2)$ steps. □

In the terminology of Computational Complexity, the previous result can be reformulated by saying that *the normalization procedure is not elementary*, i.e. it does not belong to the class $\mathcal{E}_3$ of the Grzegorczyk Hierarchy (see Odifreddi [1999], Section VIII.7).

An upper bound

Our second goal is to determine an upper bound to the number of steps needed to obtain the normal form of a term. The idea is to find a normalization procedure so simple that reduces every λ-term of length n to a normal form in at most $s(n)$ steps, for any sufficiently large n. This proves that the complexity of normalization really is at most superexponential.

Proposition 11.4.2 (Statman [1979]) *For almost all typed λ-terms, the procedure provided by the Weak Normalization Theorem* **??** *does not take more than a superexponential number of steps.*

Proof. Recall that the normalization procedure defined in 8.3.1 reduces at each step a redex $(\lambda x^\alpha . u^\beta)v^\alpha$ of greatest degree, such that in v^α no redex of greatest degree occur.

A term of length n can contain at most n redexes of greatest degree, and at most n redexes of smaller degree. The normalization procedure reduces the redexes of greatest degree one at a time, in at most n steps. At each such step, the number of redexes of smaller degree can at most be squared, as in the previous proof. After all redexes of greatest degree have been eliminated, the number of redexes of smaller degree can thus become at most

$$\underbrace{n + n^2 + (n^2)^2 + \cdots}_{n \text{ times}} \leq n^{2^n} \leq e_3(n).$$

In particular, there can be at most $e_3(n)$ redexes of a new greatest degree. Again, they are reduced one at a time, in at most $e_3(n)$ steps. At each such step, the number of redexes of smaller degree can at most be squared, and become at most

$$(e_3(n))^{2^{e_3(n)}} \leq e_3(e_3(n)) = e_3^{(2)}(n),$$

and so on. In the end, approximately $e_3^{(n)}(n)$ steps, i.e. a superexponential number, are sufficient to reduce all redexes and produce the normal form. □

In the terminology of Computational Complexity, the previous result can be reformulated by saying that *the normalization procedure is superelementary*, i.e. it belongs to the class $\mathcal{E}_4$ of the Grzegorczyk Hierarchy (see Odifreddi [1999], Section VIII.8).

An analysis of the semantical proof of the Strong Normalization given in 10.5.3 shows that the previous result actually holds for *every* normalization procedure, and not only for the particular one provided by **??**.

æ

Part D

Untyped Lambda Calculus

Chapter 12

Syntax

In Part C we studied the Typed Lambda Calculus. It would be possible to greatly enlarge its expressive power by enriching its type structure, for example by stepping to the Polymorphic Lambda Calculus. In this part we bypass all intermediate steps and perform the ultimate step toward complete freedom, by completely dropping all type restrictions.

The basic idea behind the Untyped Lambda Calculus is that *rules can express algorithms without specified domains and ranges*. For example, the rule

'give as output the input itself'

is perfectly understandable as an abstract identity function, without any particular specification of the type of arguments. But in classical mathematics it would not define a function, because its intended domain, i.e. the collection of all objects, is not a set. The Untyped Lambda Calculus is meant to be a theory of such general rules.

In set-theoretical terms, we could compare the Typed Lambda Calculus to a formal theory of *sets*, and the Untyped Lambda Calculus to a theory of *classes*. Historically, the unrestricted versions of both Set Theory and the Lambda Calculus were introduced and developed before their restricted versions. Actually, the step from the former to the latter was forced by the discovery of inconsistencies, such as Russell's Paradox. With historical insight, the Untyped Lambda Calculus is thus highly suspicious, and part of its interest lies precisely in its surprising immunity to the usual set-theoretical paradoxes.

In the following we could easily refer directly to the treatment of Chapter 8, by just dropping types everywhere. We choose a middle way, by repeating the main definitions and statements but not their motivations and proofs.

12.1 Untyped Lambda Terms

Terms

As in the Typed Lambda Calculus, terms will be seen as names for functions, their arguments and their values. The novelty here is the lack of type restrictions.

The **language** for the description of terms consists of:

- variables $x, y, \ldots$
- parentheses '(' and ')'
- dot '.'
- the term constructor λ (*lambda operator*).

This language is enough for a first approximation to a theory of untyped functions, in which we consider only unspecified atomic terms, and denote them by term *variables*. In a second approximation we can also consider specific atomic terms, representing particular objects or functions of interest, and denote them by term *constants*. The presence of term constants distinguishes an *applied* theory of functions from a *pure* one.

Definition 12.1.1 Untyped $\boldsymbol{\lambda}$-Terms (Church [1933]) *Untyped λ-terms are defined inductively as follows.*

1. **Variables**. *A variable x is a term, and x occurs free in it.*
2. **Functional Application**. *If u and v are terms, then (uv) is a term. An occurrence of a variable is free or bound in it if was so in u or v.*
3. **Functional Abstraction**. *If x is a variable and u is a term, then $(\lambda x.\, u)$ is a term. An occurrence of a variable is free or bound in it if was so in u, with the exception of the free occurrences of x in u, which become bound.*

Terms in which no variable occurs free are called **closed**.

To increase readability some parentheses can either be omitted, when no confusion arises, or written differently, e.g. as '[' and ']'. We will use the letters x, y, z, $\ldots$ for variables, and t, u, v, $\ldots$ for terms.

As in Chapter 8, we adopt the following **convention on multiple $\boldsymbol{\lambda}$-abstractions**, that defines λ-abstractions on n-tuples of variables:

$$\lambda x_1 \cdots x_n.\, u \overset{\text{def}}{=} \lambda x_1.\, (\cdots (\lambda x_n.\, u) \cdots).$$

A complementary **convention on multiple applications**, consistent with the previous one, allows us to consider the simultaneous application of a single term t to any n-tuple of terms:

$$tv_1 \cdots v_n \overset{\text{def}}{=} (\cdots (tv_1) \cdots v_n).$$

Reductions

The following rules are intuitively justified as in Chapter 8.

Definition 12.1.2 $\boldsymbol{\alpha}$-Rule. *In a given term we can change every bound occurrence of a variable with occurrences of another variable, as long as no free occurrence of any variable in any subterm of the original term becomes bound in that subterm after the change.*

Definition 12.1.3 $\boldsymbol{\beta}$-Rule. *Given terms u and v, we can step from $(\lambda x.\, u)v$ (called a* **redex***) to $u[x := v]$ (called a* **reduct**[1]*), where the latter is the result of the substitution of v for the free occurrences of x in u. We write*

$$(\lambda x.\, u)v \ \longrightarrow_{1\beta}\ u[x := v]$$

to state that one step of the β-rule has been applied to the left-hand-side to produce the right-hand-side.

Formally, $u[x := v]$ is defined by induction on u, as follows:

$$u[x := v] = \begin{cases} v & \text{if } u = x \\ u & \text{if } u = y \neq x \\ (u_1[x := v])(u_2[x := v]) & \text{if } u = u_1 u_2 \\ \lambda y.\, (u_1[x := v]) & \text{if } u = \lambda y.\, u_1, \end{cases}$$

where in the last clause we tacitly use the α-rule to ensure that the bound variable is not x itself.

The **$\boldsymbol{\alpha}$-rule** can be expressed, in terms of substitution, as follows (with the appropriate restrictions on y):

$$\lambda x.\, u = \lambda y.\, (u[x := y]).$$

As usual, a λ-term is said to be in **$\boldsymbol{\beta}$-normal form** if no application of the β-rule is possible inside it, i.e. the term does not contain any redex. We will later prove that, unlike in the Typed Lambda Calculus, *not every λ-term has a β-normal form* (12.2.1). Moreover, *not every sequence of applications of the β-rule necessarily produces the β-normal form of a term that has one* (12.2.8). However, as in the Typed Lambda Calculus, *if a β-normal form exists, then it is unique* (**Uniqueness of $\boldsymbol{\beta}$-Normal Forms**, 12.2.7).

As usual, to formally define the expression ‘to apply the β-rule *inside* a term’ we need to extend the β-rule (as defined in 12.1.3) to allow for its application not only to a term that *is* a redex, but also to any term that *contains* a redex, by inductively extending the meaning of $\longrightarrow_{1\beta}$ as follows.

[1] Some authors call it a **contractum**.

Definition 12.1.4 One-Step β-Reducibility. *The reducibility $\longrightarrow_{1\beta}$ is defined inductively by the following clauses:*

$$(\lambda x.\, u)v \ \longrightarrow_{1\beta}\ u[x := v] \tag{12.1}$$

$$\frac{u_1 \ \longrightarrow_{1\beta}\ u_2}{u_1 v \ \longrightarrow_{1\beta}\ u_2 v} \tag{12.2}$$

$$\frac{v_1 \ \longrightarrow_{1\beta}\ v_2}{u v_1 \ \longrightarrow_{1\beta}\ u v_2} \tag{12.3}$$

$$\frac{u_1 \ \longrightarrow_{1\beta}\ u_2}{\lambda x.\, u_1 \ \longrightarrow_{1\beta}\ \lambda x.\, u_2,} \tag{12.4}$$

where the first clause (i.e. the β-rule) can be thought of as an axiom, and the remaining ones as deduction rules.

A further extension, needed in the informal discussion above on normal forms, requires the application of the β-rule not only once inside a given term, but any finite number of times. This defines the notion of **β-reducibility**, which we will indicate by $\longrightarrow_{\beta}$, and of which $\longrightarrow_{1\beta}$ constitutes a single step. By definition, $\longrightarrow_{\beta}$ is simply the *reflexive and transitive closure of* $\longrightarrow_{1\beta}$. More formally:

Definition 12.1.5 β-Reducibility. *The reducibility $\longrightarrow_{\beta}$ is defined inductively by the following clauses:*

$$u \ \longrightarrow_{\beta}\ u \tag{12.1}$$

$$\frac{u_1 \ \longrightarrow_{1\beta}\ u_2}{u_1 \ \longrightarrow_{\beta}\ u_2} \tag{12.2}$$

$$\frac{u_1 \ \longrightarrow_{\beta}\ u_2 \qquad u_2 \ \longrightarrow_{\beta}\ u_3}{u_1 \ \longrightarrow_{\beta}\ u_3} \tag{12.3}$$

where the first clause can be thought of as an axiom, and the remaining ones as deduction rules.

Equality

The final extension of β-reducibility that we consider is the notion of **β-equality**, which we will indicate by $=_{\beta}$, and is defined as the *symmetric and transitive closure of* $\longrightarrow_{\beta}$. More formally:

Definition 12.1.6 β-Equality. *The equivalence relation $=_{\beta}$ is defined inductively by the following clauses:*

$$\frac{u_1 \ \longrightarrow_{\beta}\ u_2}{u_1 \ =_{\beta}\ u_2} \tag{12.4}$$

$$\frac{u_1 \ =_\beta \ u_2}{u_2 \ =_\beta \ u_1} \tag{12.5}$$

$$\frac{u_1 =_\beta u_2 \qquad u_2 =_\beta u_3}{u_1 =_\beta u_3,} \tag{12.6}$$

where the first clause can be thought of as an axiom, and the remaining ones as deduction rules.

The next property is proved as in 8.1.7 and 8.5.8, using the untyped version of the Diamond Property (12.2.6).

Proposition 12.1.7 *Two terms are β-equal if and only if they reduce to a common term.*

In the Untyped Lambda Calculus the notion of β-equality cannot be characterized solely in terms of normal forms, as it was the case for the Typed Lambda Calculus (see 8.4.6). The connections between β-equality and normal forms are determined in the next exercises.

Exercises 12.1.8 a) *Any term β-equal to a term having a normal form, also has a normal form.* (Hint: by 12.2.6 and 8.5.8.)

b) *Terms having a normal form are β-equal if and only if they have the same normal form.*

c) *Not all terms without a normal form are β-equal.*(Hint: the terms $\Delta\Delta$ and $\Delta_y\Delta_y$ used in the proof of 12.2.2 are not β-equal.)

Combinators ⋆

As for the Typed Lambda Calculus, we can present the Untyped Lambda Calculus in a *synthetical* way, in which a few λ-terms called **atomic combinators** are selected, and **combinators** are built up from them by means of application alone. The question again arises of finding nontrivial atomic combinators, such that the terms built up from them and the variables by means of application alone, which we will call **combinatorial terms**,[2] somehow represent all λ-terms. The answer is given by the next result, whose proof is like the one in 8.2.1.

Theorem 12.1.9 Functional Completeness (Schönfinkel [1924], Curry [1930]) *Define the* **atomic combinators** *as follows:*

1. $\mathbf{I} = \lambda x.\, x$
2. $\mathbf{K} = \lambda yx.\, y$

[2] The usual combinators are *closed* combinatorial terms, i.e. the ones without free variables.

3. $\mathbf{S} = \lambda xyz.\,(xz)(yz)$.

Then for every λ*-term* t *there is a* **combinatorial term** t_c *built up from atomic combinators and variables by application alone, and such that* t_c *is* β*-reducible to* t.

The main difference with the Typed Lambda Calculus is that we only have here three combinators, as opposed to three *families* of them, one for each possible type. As usual, **I** is derivable from the other two as **SKK**, and thus only two combinators are enough to synthesize every λ-term, by composition alone and starting from the variables.

Having shown how combinators are actually sufficient to define all λ-terms, we can take a last step and develop a **theory of combinators** independently of the Untyped Lambda Calculus, and as an alternative approach to it.

The notion of a *combinator* is defined inductively, as in 12.1.1, using in a first approximation only the two constants **K** and **S** (in a second approximation, the presence of other combinator constants distinguishes the *pure* from an *applied* theory of combinators):

1. The constants **K** and **S** are combinators.
2. If **C** and **D** are combinators, then (**CD**) is a combinator.

To increase readability some parentheses can be omitted, when no confusion arises. We will use the letters x, y, z, … for variables, and **C**, **D**, **E**, … for combinators.

The notion of *combinatorial* β*-reducibility* $\longrightarrow_{c\beta}$ is defined in analogy with β-reducibility, by replacing the β-rule with its two instances needed to give the constants **K** and **S** the appropriate operational meaning. Precisely, $\longrightarrow_{c\beta}$ is the reflexive and transitive closure of the single step reducibility $\longrightarrow_{1c\beta}$, defined inductively as:

$$\mathbf{KCD} \longrightarrow_{1c\beta} \mathbf{C}$$

$$\mathbf{SCDE} \longrightarrow_{1c\beta} (\mathbf{CE})(\mathbf{DE})$$

$$\frac{\mathbf{C}_1 \longrightarrow_{1c\beta} \mathbf{C}_2}{\mathbf{C}_1\mathbf{D} \longrightarrow_{1c\beta} \mathbf{C}_2\mathbf{D}}$$

$$\frac{\mathbf{D}_1 \longrightarrow_{1c\beta} \mathbf{D}_2}{\mathbf{C}\mathbf{D}_1 \longrightarrow_{1c\beta} \mathbf{C}\mathbf{D}_2.}$$

The notion of *combinatorial* β*-equality* $=_{c\beta}$ is defined as the symmetric and transitive closure of $\longrightarrow_{c\beta}$.

A combinator is in β*-normal form* if no reduction as above is possible inside it or, equivalently, if it does not contain any subcombinator of the form **KCD** or **SCDE**.

As for the Typed Theory of Combinators and the Typed Lambda Calculus, it is possible to define mutual translations of a combinator $\mathbf{C}$ into a λ-term $\mathbf{C}_\lambda$, and of a λ-term t into a combinator t_c (see 8.6.6 and 8.6.7). In the presence of extensionality rules the two translations turn out to be inverse one of the other, i.e.

$$(\mathbf{C}_\lambda)_c = \mathbf{C} \qquad \text{and} \qquad (t_c)_\lambda = t$$

(see 8.6.9 and 8.6.10).

12.2 Normal Forms

We consider in this section the topic of normal forms. It is here that the Typed and Untyped Lambda Calculi radically depart way. In the former, every term has a unique normal form, and any reduction procedure eventually produces it. In the latter, *not every term has a normal form, and even when it does, not every reduction procedure eventually produces it*.

However, one result that was true of the Typed Lambda Calculus remains true, namely that *normal forms are unique, when they exist*. This shows that the Untyped Lambda Calculus can still be considered as a calculus of functions, although not anymore of total ones.

Terms without normal form

The next result shows that even the Weak Normalization Theorem fails for the Untyped Lambda Calculus.

Proposition 12.2.1 *There is a term without normal form.*

Proof. The idea is to construct a term that perpetually selfreproduces, and hence it never reaches a normal form. To be able to selfreproduce, such a term must be (or at least contain) an application, otherwise we could not apply the β-rule to it (or a subterm of it). We are thus looking for a term of the form uv, where u is a λ-abstraction. Moreover, for uv to be able to reproduce itself, u must produce an application when applied to v. The simplest term u that is an abstraction and produces an application is $\Delta = \lambda x.\, xx$. Then

$$\Delta v \longrightarrow_\beta vv.$$

By applying Δ to itself, i.e. by letting $v = \Delta$, we get

$$\Delta\Delta \longrightarrow_\beta \Delta\Delta,$$

and thus $\Delta\Delta$ selfreproduces.

Of the two occurrences of λ in $\Delta\Delta = (\lambda x.\, xx)(\lambda x.\, xx)$, the rightmost one cannot be reduced because it is not applied to any term. Thus the only possible application of the β-rule inside $\Delta\Delta$ is precisely the one that forces it to selfreproduce, and no normal form exists. □

Notice that the selfreproduction of $\Delta\Delta$ occurs in exactly one step of β-reduction. The proof just given actually shows that $\Delta\Delta$ is the only term for which this happens.

Fixed points

The technique just used can be easily extended to obtain the following crucial result, for which we provide two different proofs.

Theorem 12.2.2 Fixed Point Combinator (Kleene [1936], Turing [1937], Curry [1942], Rosenbloom [1950]) *Every untyped λ-term has a* **fixed point***, i.e. for every u there is v such that*

$$v \longrightarrow_\beta uv.$$

Actually, there is a **fixed point combinator** *$\mathcal{Y}$ that, when applied to u, produces a fixed point of it:*

$$\mathcal{Y}u =_\beta u(\mathcal{Y}u).$$

First Proof. We start with an informal argument. Recall that in 12.2.1 we defined

$$\Delta = \lambda x.\, xx,$$

and noticed that

$$\Delta v \longrightarrow_\beta vv,$$

so that

$$\Delta\Delta \longrightarrow_\beta \Delta\Delta.$$

We now want a term that reproduces not itself, but u applied to itself. The natural guess is to use

$$\Delta_u = \lambda x.\, u(xx).$$

Indeed,

$$\Delta_u v \longrightarrow_\beta u(vv),$$

and

$$\Delta_u\Delta_u \longrightarrow_\beta u(\Delta_u\Delta_u).$$

Actually, $\Delta_u\Delta_u$ is obtained uniformly in u, and by abstracting u we obtain the required fixed point combinator:

$$\mathcal{Y} = \lambda y.\, \Delta_y\Delta_y = \lambda y.\, (\lambda x.\, y(xx))(\lambda x.\, y(xx)).$$

Then

$$\begin{array}{rcl} \mathcal{Y}u & = & (\lambda y.\, \Delta_y \Delta_y)u \\ & \rightarrow_\beta & \Delta_u \Delta_u \\ & \rightarrow_\beta & u(\Delta_u \Delta_u) \\ & =_\beta & u(\mathcal{Y}u). \end{array}$$

Notice that the last step is really only a β-equality, and not a β-reduction, because we are going from the reduct $\Delta_u \Delta_u$ to the redex $(\lambda y.\, \Delta_y \Delta_y)u = \mathcal{Y}u$, i.e. in the direction opposite to $\longrightarrow_\beta$.

Second Proof. We now give a proof based on the contrapositive of the diagonal method, in the style of Owings [1973]. Given a term u, suppose $v \longrightarrow_\beta uv$ as required. Then v is β-equal to a term of the form $t_i t_j$, for some t_i and t_j. We can thus think of u as acting on terms of the form $t_i t_j$. Consider then an enumeration of all pairs of λ-terms:

$$\begin{array}{lllll} t_0t_0 & t_0t_1 & t_0t_2 & t_0t_3 & \cdots \\ t_1t_0 & t_1t_1 & t_1t_2 & \cdots & \\ t_2t_0 & t_2t_1 & t_2t_2 & \cdots & \\ t_3t_0 & \cdots & \cdots & \cdots & \\ \cdots & & & & \end{array}$$

and the effect of u on the diagonal:

$$u(t_0t_0) \quad u(t_1t_1) \quad u(t_2t_2) \quad \cdots$$

This is β-equal to a row of the previous matrix. More precisely, to the n-th one, with n such that

$$t_n = \lambda x.\, u(xx),$$

because

$$t_n t_i \longrightarrow_\beta u(t_i t_i).$$

In particular,

$$t_n t_n \longrightarrow_\beta u(t_n t_n),$$

i.e. $t_n t_n$ is a fixed point of u. Actually,

$$t_n t_n = (\lambda x.\, u(xx))(\lambda x.\, u(xx))$$

is obtained uniformly in u, and by abstracting u we obtain the required fixed point combinator:

$$\mathcal{Y} = \lambda y.\, (\lambda x.\, y(xx))(\lambda x.\, y(xx)).$$

At this point, we can proceed as in the previous proof. □

$\mathcal{Y}$ is sometimes called the **paradoxical combinator**, because it embodies the argument used in *Russell's paradox*. The connection is established by the following correspondence:

Set Theory	Lambda Calculus
element	argument
set	function
membership	application
set formation	λ-abstraction.

Russell's paradox is obtained by considering the set

$$A = \{x : x \notin x\}.$$

Then

$$x \in A \;\Leftrightarrow\; x \notin x,$$

and

$$A \in A \;\Leftrightarrow\; A \notin A.$$

The last assertion is a contradiction.

In terms of the Lambda Calculus, negation can be considered as term u that is never the identity. Since membership corresponds to application, and set formation to λ-abstraction, the set A corresponds to the term

$$\Delta_u = \lambda x.\, u(xx).$$

Then

$$\Delta_u x =_\beta u(xx),$$

and

$$\Delta_u \Delta_u =_\beta u(\Delta_u \Delta_u).$$

The last assertion is now true, and not a contradiction. Rather, from it we deduce that a term u that is never the identity cannot exist.

Uniqueness of normal forms

In Chapter 8 we have provided two different proofs of uniqueness of normal forms. One (8.4.5) used the Strong Normalization Theorem and obviously has no counterpart here, since even the Weak Normalization Theorem fails. The other (8.5.7) followed from the Church-Rosser Theorem (8.5.6), and goes through here with no change, by simply erasing types.

We only restate the relevant definitions and results, and refer to Section 8.4 for motivations and proofs.

Definition 12.2.3 Parallel Reducibility (Tait, Martin-Löf) *The reducibility $\Longrightarrow$ is defined inductively by the following clauses:*

$$u \Longrightarrow u \tag{12.1}$$

$$\frac{u \Longrightarrow u_1 \quad v \Longrightarrow v_1}{uv \Longrightarrow u_1 v_1} \tag{12.2}$$

$$\frac{u \Longrightarrow u_1}{\lambda x.\, u \Longrightarrow \lambda x.\, u_1} \tag{12.3}$$

$$\frac{u \Longrightarrow u_1 \quad v \Longrightarrow v_1}{(\lambda x.\, u)v \Longrightarrow u_1[x := v_1],} \tag{12.4}$$

where the first clause can be thought of as an axiom, and the remaining one as deduction rules.

As a special case of equation 12.4 we have

$$\frac{u \Longrightarrow u \quad v \Longrightarrow v}{(\lambda x.\, u)v \Longrightarrow u[x := v],}$$

where the top line consists of axioms. Thus $\longrightarrow_{1\beta}$ implies $\Longrightarrow$. Inductively, we easily see that $\Longrightarrow$ implies $\longrightarrow_\beta$. Then $\longrightarrow_\beta$ *is the transitive closure of* $\Longrightarrow$, since it is the transitive closure of $\longrightarrow_{1\beta}$.

The next results are proved as in 8.5.4, 8.5.5, 8.5.6 and 8.5.7.

Proposition 12.2.4 Strong Diamond Property for $\Longrightarrow$ (Takahashi [1995]) *For any term t there is a term t^* such that if $t \Longrightarrow t_1$, then $t_1 \Longrightarrow t^*$.*

Corollary 12.2.5 Diamond Property for $\Longrightarrow$ (Tait, Martin-Löf) *If t_1 and t_2 are terms obtained from t by $\Longrightarrow$, then there is a term t^* which can be obtained from t_1 and t_2 by $\Longrightarrow$. Graphically,*

$$\begin{tikzcd} t \arrow[r, Rightarrow] \arrow[d, Rightarrow] & t_1 \arrow[d, Rightarrow] \\ t_2 \arrow[r, Rightarrow] & t^*. \end{tikzcd}$$

Theorem 12.2.6 Diamond Property for $\longrightarrow_\beta$ (Church and Rosser [1936]) *If t_1 and t_2 are terms obtained from t by $\longrightarrow_\beta$, then there is a term t^* which can be obtained from t_1 and t_2 by $\longrightarrow_\beta$. Graphically,*

$$\begin{tikzcd} t \arrow[r] \arrow[d] & t_1 \arrow[d] \\ t_2 \arrow[r] & t^*. \end{tikzcd}$$

Corollary 12.2.7 Uniqueness of Normal Forms. *Every untyped λ-term has at most one normal form.*

Normalizing reduction strategies

Due to the presence of selfreproducing terms and the possibility of cancellations, a term may have a normal form, because bad subterms are eventually cancelled out, but some reduction strategy may not produce it, because it may get stuck in the evaluation of a subterm without normal form.

Proposition 12.2.8 *There is a term with a normal form, and such that some reduction strategy does not produce it.*

Proof. Consider $(\lambda xy.y)(\Delta\Delta)$. Then:

- by reducing the outermost λ we get $\lambda y.\, y$, which is in normal form
- by always reducing the outermost λ in $\Delta\Delta$, we continue to obtain the term itself, since $\Delta\Delta$ selfreproduces. □

The problem then arises of whether there is a reduction strategy that would always produce the normal form of a term, whenever it exists. We get a positive answer by using the following notions.

Definition 12.2.9 *A one-step β-reduction $u \longrightarrow_{1\beta} v$ is called:*

1. *a* **head reduction** *($u \longrightarrow_{1h} v$) if*

$$u = \lambda y_1 \cdots y_n.\, (\lambda x.\, u_0)u_1 \cdots u_m$$

and v is obtained from it by reducing $(\lambda x.\, u_0)u_1$, i.e.

$$v = \lambda y_1 \cdots y_n.\, (u_0[x := u_1])u_2 \cdots u_m.$$

2. *an* **internal reduction** *($u \longrightarrow_{1i} v$) if it is not a head reduction, i.e.*

$$u = \lambda y_1 \cdots y_n.\, (\lambda x.\, u_0)u_1 \cdots u_m \qquad \textit{or} \qquad \lambda y_1 \cdots y_n.\, zu_0 \cdots u_m,$$

and the reduction is performed inside one of the terms u_i.

3. *a* **leftmost reduction** *($u \longrightarrow_{1l} v$) if it reduces the leftmost λ that can be reduced.*

We define $\longrightarrow_h$, $\longrightarrow_i$ and $\longrightarrow_l$ as in 12.1.5, i.e. as the reflexive and transitive closure of $\longrightarrow_{1h}$, $\longrightarrow_{1i}$ and $\longrightarrow_{1l}$.

We define $\Longrightarrow_i$ as in 12.2.3, i.e. as the parallel version of $\longrightarrow_{1i}$.

Obviously, there is no parallel version of either $\longrightarrow_{1h}$ or $\longrightarrow_{1l}$, because at most one head or leftmost reduction can be performed on a given term. It is instead possible to perform different internal reductions, and $\Longrightarrow_i$ allows us to perform some of them in parallel.

We now prove that a parallel reduction can always be factored into a sequence of head reductions, followed by an internal parallel reduction.

Theorem 12.2.10 Normal Form of Parallel Reductions (Takahashi [1989]) *If $u \Longrightarrow v$, then there exists a term t such that $u \longrightarrow_h t \Longrightarrow_i v$.*

Proof. We prove by induction on $u \Longrightarrow v$ that there exist terms $t_i \Longrightarrow v$ such that

$$u \longrightarrow_{1h} t_1 \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_n \Longrightarrow_i v.$$

1. $u \Longrightarrow u$
   In this case there is nothing to prove.

2. $u_1u_2 \Longrightarrow v_1v_2$, with $u_1 \Longrightarrow v_1$ and $u_2 \Longrightarrow v_2$
   By the induction hypothesis on $u_1 \Longrightarrow v_1$, there exist terms $t_i \Longrightarrow v_1$ such that
   $$u_1 \longrightarrow_{1h} t_1 \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_n \Longrightarrow_i v_1.$$
   We append u_2 everywhere, and notice that:

   (a) If none of $u_1, t_1, \ldots, t_n$ is a λ-abstraction, then
   $$u_1u_2 \longrightarrow_{1h} t_1u_2 \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_nu_2$$
   because no new head reduction becomes possible by appending u_2.

   (b) Since $t_n \Longrightarrow_i v_1$ and $u_2 \Longrightarrow v_2$, then
   $$t_nu_2 \Longrightarrow_i v_1v_2.$$
   Indeed, any reduction internal in t_n remains internal in t_nu_2. And any reduction in u_2 is internal in t_nu_2.

   Then
   $$u_1u_2 \longrightarrow_{1h} t_1u_2 \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_nu_2 \Longrightarrow_i v_1v_2,$$
   i.e. Case 2 holds if none of $u_1, t_1, \ldots, t_n$ is a λ-abstraction. Otherwise, there are two cases.

   (c) If u_1 is a λ-abstraction, then
   $$u_1u_2 \Longrightarrow_i v_1v_2.$$
   Indeed, any reduction in u_1 is internal in u_1u_2, because u_1 is a λ-abstraction. And any reduction in u_2 is internal in u_1u_2. Then Case 2 holds trivially.

(d) Otherwise, let t_m be the first such term which is a λ-abstraction. Then

$$u_1u_2 \longrightarrow_{1h} t_1u_2 \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_mu_2 \Longrightarrow_i v_1v_2.$$

Indeed, the head reductions hold as in (a), because none of $u_1, t_1, \ldots, t_{m-1}$ is a λ-abstraction. And the internal parallel reduction holds as in (c), because $t_m \Longrightarrow v_1$ by the additional induction hypothesis, which was introduced precisely to make this case work.

3. $\lambda x.\, u \Longrightarrow \lambda x.\, v$, with $u \Longrightarrow v$
By the induction hypothesis on $u \Longrightarrow v$, there exist terms $t_i \Longrightarrow v$ such that

$$u \longrightarrow_{1h} t_1 \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_n \Longrightarrow_i v.$$

If we prefix λx everywhere, we trivially get $\lambda x.\, t_i \Longrightarrow \lambda x.\, v$ and

$$\lambda x.\, u \longrightarrow_{1h} \lambda x.\, t_1 \longrightarrow_{1h} \cdots \longrightarrow_{1h} \lambda x.\, t_n \Longrightarrow_i \lambda x.\, v.$$

4. $(\lambda y.\, u_1)v_1 \Longrightarrow u_2[y := v_2]$, with $u_1 \Longrightarrow u_2$ and $v_1 \Longrightarrow v_2$
By the inductive hypothesis on $u_1 \Longrightarrow u_2$, there exist terms $t_i \Longrightarrow u_2$ such that

$$u_1 \longrightarrow_{1h} t_1 \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_n \Longrightarrow_i u_2.$$

If we substitute v_1 for y everywhere, we get

$$(\lambda y.\, u_1)v_1 \longrightarrow_{1h} u_1[y := v_1] \longrightarrow_{1h} t_1[y := v_1] \longrightarrow_{1h} \cdots \longrightarrow_{1h} t_n[y := v_1].$$

Notice that:

- $u_1[y := v_1] \Longrightarrow u_2[y := v_2]$
This follows from 8.5.3, by the hypotheses $u_1 \Longrightarrow u_2$ and $v_1 \Longrightarrow v_2$.
- $t_i[y := v_1] \Longrightarrow u_2[y := v_2]$
This follows from 8.5.3, by the inductive hypothesis $t_i \Longrightarrow u_2$ and the hypothesis $v_1 \Longrightarrow v_2$.

To conclude the proof, it would thus be enough to prove

$$t_n[y := v_1] \Longrightarrow_i u_2[y := v_2].$$

We do this below, but in one case we can only prove that

$$t_n[y := v_1] \longrightarrow_h \Longrightarrow_i u_2[y := v_2].$$

This is still sufficient, but it introduces additional terms between $t_n[y := v_1]$ and $u_2[y := v_2]$, which must be proved to be reducible to $u_2[y := v_2]$ by means of parallel reductions.

By the induction hypothesis on $v_1 \Longrightarrow v_2$, we have $v_1 \longrightarrow_h \Longrightarrow_i v_2$, with intermediate terms reducible to v_2 by parallel reductions. We now prove that

$$\frac{t_n \Longrightarrow_i u_2 \qquad v_1 \longrightarrow_h \Longrightarrow_i v_2}{t_n[y := v_1] \longrightarrow_h \Longrightarrow_i u_2[y := v_2].}$$

Since we step from t_n to u_2 by internal reductions, there are two cases:

(a) $t_n = \lambda x_1 \cdots x_p .\, z s_1 \cdots s_q$ and $u_2 = \lambda x_1 \cdots x_p .\, z s_1^* \cdots s_q^*$, where $s_i \Longrightarrow s_i^*$. Since the λ's can be added later on, as in Case 3, we can restrict attention to the rest of the terms. There are two subcases:

- if $z = y$, then

$$\begin{array}{rcl}(z s_1 \cdots s_q)[y := v_1] & = & v_1(s_1[y := v_1]) \cdots (s_q[y := v_1]) \\ & \longrightarrow_h \Longrightarrow_i & v_2(s_1^*[y := v_2]) \cdots (s_q^*[y := v_2]) \\ & = & (z s_1^* \cdots s_q^*)[y := v_2]\end{array}$$

  by $v_1 \longrightarrow_h \Longrightarrow_i v_2$ and $s_i \Longrightarrow s_i^*$.
  The condition that the additional terms thus introduced are reducible to $u_2[y := v_2]$ by parallel reductions, follows from the induction hypothesis on $v_1 \longrightarrow_h \Longrightarrow_i v_2$ and from 8.5.3, which implies

$$\frac{s_i \Longrightarrow s_i^* \qquad v_1 \Longrightarrow v_2}{s_i[y := v_1] \Longrightarrow s_i^*[y := v_2].}$$

- if $z \neq y$, then

$$\begin{array}{rcl}(z s_1 \cdots s_q)[y := v_1] & = & z(s_1[y := v_1]) \cdots (s_q[y := v_1]) \\ & \Longrightarrow_i & z(s_1^*[y := v_2]) \cdots (s_q^*[y := v_2]) \\ & = & (z s_1^* \cdots s_q^*)[y := v_2]\end{array}$$

  by $s_i \Longrightarrow s_i^*$ and $v_1 \Longrightarrow v_2$.

(b) $t_n = \lambda x_1 \cdots x_p .\, (\lambda z .\, s_0) s_1 \cdots s_q$ and $u_2 = \lambda x_1 \cdots x_p .\, (\lambda z .\, s_0^*) s_1^* \cdots s_q^*$. Since we may suppose $z \neq y$ by the α-rule, we can proceed as in the second subcase of (a) above. □

Proposition 12.2.11 First Invertibility Property (Takahashi [1989]) *An internal parallel reduction followed by head reductions can be inverted, in the sense of being replaced by head reductions followed by an internal parallel reduction. Schematically:*

$$\frac{u \Longrightarrow_i \longrightarrow_h v}{u \longrightarrow_h \Longrightarrow_i v.}$$

Proof. By induction on the number of head reductions, it is enough to restrict attention to the case $u \Longrightarrow_i t \longrightarrow_{1h} v$. Then

$$\begin{array}{rcl} u & = & \lambda y_1 \cdots y_n.\, (\lambda x.\, u_0)u_1u_2 \cdots u_m \\ t & = & \lambda y_1 \cdots y_n.\, (\lambda x.\, t_0)t_1t_2 \cdots t_m \\ v & = & \lambda y_1 \cdots y_n.\, (t_0[x := t_1])t_2 \cdots t_m, \end{array}$$

where $u_i \Longrightarrow t_i$, although not necessarily $u_i \Longrightarrow_i t_i$. If we let

$$s = \lambda y_1 \cdots y_n.\, u_0[x := u_1]u_2 \cdots u_m,$$

then $u \longrightarrow_{1h} s \Longrightarrow v$. By 12.2.10, $s \Longrightarrow v$ can be factored into a head reduction and an internal parallel reduction, i.e. there exists s_1 such that $s \longrightarrow_h s_1 \Longrightarrow_i v$. Then $u \longrightarrow_h s_1 \Longrightarrow_i v$. Schematically:

$$\begin{array}{c} u \Longrightarrow_i t \longrightarrow_{1h} v \\ \hline u \longrightarrow_{1h} s \Longrightarrow v \\ \hline u \longrightarrow_{1h} s \longrightarrow_h s_1 \Longrightarrow_i v \\ \hline u \longrightarrow_h s_1 \Longrightarrow_i v. \quad \square \end{array}$$

We now prove that a β-reduction can always be factored into a sequence of head reductions, followed by a sequence of internal reductions.

Corollary 12.2.12 Normal Form of Reductions (Mitschke [1979]) *If $u \longrightarrow_\beta v$, then there exists a term t such that $u \longrightarrow_h t \longrightarrow_i v$.*

Proof. A β-reduction can be factored into a sequence of 1β-reductions, and hence of parallel reductions. By 12.2.10, each parallel reduction can be factored into a sequence of head reductions, and an internal parallel reduction. By the invertibility noticed above, this can be turned into a sequence of head reductions, followed by a sequence of internal (parallel) reductions. Schematically:

$$\begin{array}{c} u \longrightarrow_\beta v \\ \hline u \longrightarrow_{1\beta} \cdots \longrightarrow_{1\beta} v \\ \hline u \Longrightarrow \cdots \Longrightarrow v \\ \hline u \longrightarrow_h \Longrightarrow_i \cdots \longrightarrow_h \Longrightarrow_i v \\ \hline u \longrightarrow_h \Longrightarrow_i \cdots \Longrightarrow_i v \\ \hline u \longrightarrow_h \longrightarrow_i \cdots \longrightarrow_i v \\ \hline u \longrightarrow_h \longrightarrow_i v. \end{array}$$

Notice that the next to last step from $\Longrightarrow_i$ to $\longrightarrow_i$ is needed because (internal) parallel reductions are not transitive. $\square$

A term is said to be in *head normal form* when it cannot be reduced by means of head reductions. The previous corollary implies that if a term has a head normal form, then the latter can be produced by a sequence of head reductions.

With a little more effort we can prove that if a term has a *normal form*, then the latter can be produced by a sequence of leftmost reductions. In other words, leftmost reduction is the normalization strategy we were looking for.

Corollary 12.2.13 Leftmost Reduction Strategy (Church and Rosser [1936]) *If $u \longrightarrow_\beta v$ and v is in normal form, then $u \longrightarrow_l v$.*

Proof. We proceed by induction on $u \longrightarrow_\beta v$. By 12.2.12, there exists a term t such that $u \longrightarrow_h t \longrightarrow_i v$. Since v is in normal form, there exist terms $v_1, \ldots, v_m$ in normal form, such that

$$v = \lambda y_1 \cdots y_n.\, z v_1 \cdots v_m.$$

Since v is obtained from t by internal reductions,

$$t = \lambda y_1 \cdots y_n.\, z t_1 \cdots t_m,$$

with $t_i \longrightarrow_\beta v_i$. By the induction hypothesis, $t_i \longrightarrow_l v_i$. Then

$$\begin{array}{rll} u & \longrightarrow_h & \lambda y_1 \cdots y_n.\, z t_1 t_2 \cdots t_m \\ & \longrightarrow_l & \lambda y_1 \cdots y_n.\, z v_1 t_2 \cdots t_m \\ & \longrightarrow_l & \lambda y_1 \cdots y_n.\, z v_1 v_2 \cdots t_m \\ & \longrightarrow_l & \cdots \\ & \longrightarrow_l & \lambda y_1 \cdots y_n.\, z v_1 v_2 \cdots v_m. \end{array}$$

Since a head reduction is a leftmost reduction, $u \longrightarrow_l v$ □

Exercises 12.2.14 a) **Quasi-Head Reduction Theorem.** *In a sequence of β-reductions starting with a term having head normal form, there cannot be infinitely many head reductions. It follows that if a term has a head normal form, then the latter can be produced by any reduction strategy that uses head reductions unboundedly often.* (Takahashi [1995]) (Hint: by 12.2.12 and invertibility, from a sequence of β-reductions containing infinitely many head reductions we can extract arbitrarily long finite sequences of head reductions starting with the initial term. Then such a term cannot have a head normal form.)

b) **Quasi-Leftmost Reduction Theorem.** *In a sequence of β-reductions starting with a term having normal form, there cannot be infinitely many leftmost reductions. It follows that if a term has a normal form, then the latter can be produced by any reduction strategy that uses leftmost reductions unboundedly often.* (Barendregt [1981]) (Hint: by induction on the first term u of the sequence. Suppose u has a normal form $\lambda y_1 \cdots y_n.\, z u_1 \cdots u_m$. By part a), in a sequence of β-reductions starting with u there cannot be infinitely many head reductions. If there are infinitely many leftmost reductions, there must be a leftmost but not head reduction starting with a term v in the sequence. Then v is in head normal form, and by the Church-Rosser Theorem it must be of the form $\lambda y_1 \cdots y_n . z v_1 \cdots v_m$, with $v_i \longrightarrow_\beta u_i$. By the induction hypothesis, there can be only finitely many leftmost reductions in any sequence of β-reductions starting with any v_i, and hence only finitely many in any sequence starting with v.)

c) **Standardization Theorem.** *If $u \longrightarrow_\beta v$, it is possible to go from u to v by a sequence of reductions in which the contracted redexes move from left to right. In other words, once a redex is reduced, all redexes to its left become frozen and cannot be reduced anymore.* (Curry and Feys [1958]) (H nduction on u, using 12.2.12.)

Notice that our treatment reverses the historically order. The Leftmost Reduction Theorem was originally proved in Church and Rosser [1936] by a complicated proof. Successive simplifications were given first by Curry and Feys [1958], using the Standardization Theorem, and then by Mitschke [1979], using the Normal Form Theorem for Reductions. Finally, Takahashi [1989] gave the proof above, using the Normal Form Theorem for Parallel Reductions. On their turn, parallel reductions were first introduced by Tait and Martin-Löf to give a simplified proof of the Church-Rosser Theorem, which again was originally proved by Church and Rosser [1936] by a complicated proof.

12.3 Extensionality

As for the Typed Lambda Calculus, we now add a rule that identifies intensionally different terms computing the same function, and study the extensional version of the Untyped Lambda Calculus thus obtained. Since much of the work done in Section 8.6 extends without changes, we concentrate on the novel aspects.

As usual, we will explicitly refer to notions, symbols and results of the previous sections by attaching a β to them. For example, we will talk of β-normal forms, parallel β-reductions $\Rightarrow_\beta$ and β-normalizing reduction strategies.

Extensional reductions

We do not repeating the discussion of Section 8.6 on the various equivalent extensionality rules, and consider only the following.

Definition 12.3.1 $\boldsymbol{\eta}$-rule. *Given a term u and a variable x not occurring free in it, we can step from $\lambda x . ux$ (called an **$\boldsymbol{\eta}$-redex**) to u (called an **$\boldsymbol{\eta}$-reduct**). We write*

$$\lambda x . ux \longrightarrow_{1\eta} u$$

to state that one step of the η-rule has been applied to the left hand side to produce the right hand side.

The following definition is the analogue of 12.1.4, 12.1.5 and 12.1.6.

Definition 12.3.2 η-Reducibility and η-Equality. *The reducibility $\longrightarrow_{1\eta}$ is defined inductively, by replacing in 12.1.4 the first clause by the following:*

$$\lambda x.\, ux \longrightarrow_{1\eta} u,$$

when x is not free in u.

The reducibility $\longrightarrow_{\eta}$ is the reflexive and transitive closure of $\longrightarrow_{1\eta}$.

The relation $=_{\eta}$ is the symmetric and transitive closure of $\longrightarrow_{\eta}$.

We say that a term is in **η-normal form** if no application of the η-rule is possible in it. Then, *given a term, any sequence of applications of η-reductions to it will eventually produce a term in η-normal form* (**Strong η-Normalization**), *that is independent of the chosen sequence of η-reductions* (**Uniqueness of η-Normal Forms**). The proofs are the same as for the Typed Lambda Calculus. In particular, it is still true that any application of the η-rule decreases the length of the term to which it is applied.

The next definition considers the η-rule not as a replacement of the β-rule, as above, but as a supplement to it.

Definition 12.3.3 $\beta\eta$-Reducibility and $\beta\eta$-Equality. *The reducibility $\longrightarrow_{1\beta\eta}$ is defined inductively, by adding to 12.1.4 the following clause:*

$$\lambda x.\, ux \longrightarrow_{1\eta} u,$$

when x is not free in u.

The reducibility $\longrightarrow_{\beta\eta}$ is the reflexive and transitive closure of $\longrightarrow_{1\beta\eta}$.

The relation $=_{\beta\eta}$ is the symmetric and transitive closure of $\longrightarrow_{\beta\eta}$.

As above, we say that a term is in **$\beta\eta$-normal form** if no application of the β-rule or η-rule is possible in it. By 12.2.1, *the $\beta\eta$-normal form of an untyped term does not always exist*. However, *if a $\beta\eta$-normal form exists, then it is unique* (**Uniqueness of $\beta\eta$-normal Forms**, 12.3.8) *and it can be obtained by leftmost $\beta\eta$-reductions* (**Leftmost $\beta\eta$-Reduction Strategy**, 12.3.13).

Normal Forms

In Chapter 8 we have provided two different proofs of the Diamond Property for $\longrightarrow_{\beta\eta}$, and hence of $\beta\eta$-normal forms. One proof (8.6.17) used a class of terms defined by induction on types and obviously has no counterpart here. The other proof (8.6.15) used parallel $\beta\eta$-reductions and goes through here with no changes, by simply erasing the types.

We only restate the relevant definitions and results, and refer to Section 8.6 for proofs. The next concept is not needed for the Diamond Property, but it will be useful later on.

Definition 12.3.4 Parallel η-Reducibility (Tait, Martin-Löf) *The reducibility $\Longrightarrow_\eta$ is defined inductively, by replacing in 12.2.3 the first clause by the following:*

$$\frac{u \Longrightarrow_\eta u_1}{\lambda x.\, ux \Longrightarrow_\eta u_1,} \tag{12.5}$$

when x is not free in u.

As a special case of equation 12.5 we have

$$\frac{u \Longrightarrow_\eta u}{\lambda x.\, ux \Longrightarrow_\eta u,}$$

where the top line consists of an axiom. Thus, as usual, $\longrightarrow_\eta$ *is the transitive closure of* $\Longrightarrow_\eta$.

Definition 12.3.5 Parallel $\beta\eta$-Reducibility (Tait, Martin-Löf) *The reducibility $\Longrightarrow_{\beta\eta}$ is defined inductively, by adding to 12.2.3 the following clause:*

$$\frac{u \Longrightarrow_{\beta\eta} u_1}{\lambda x.\, ux \Longrightarrow_{\beta\eta} u_1,} \tag{12.6}$$

when x is not free in u.

Once again, as usual, $\longrightarrow_{\beta\eta}$ *is the transitive closure of* $\Longrightarrow_{\beta\eta}$. We can now prove the next results as in 8.6.14, **??** and 8.6.16.

Proposition 12.3.6 (Takahashi [1995]) *For any term t there is a term t^* such that if $t \Longrightarrow_{\beta\eta} t_1$, then $t_1 \Longrightarrow_{\beta\eta} t^*$.*

Theorem 12.3.7 Diamond Property for $\beta\eta$-Reduction (Curry and Feys [1958]) *If t_1 and t_2 are terms obtained from t by $\longrightarrow_{\beta\eta}$, then there is a term t^* which can be obtained from t_1 and t_2 by $\longrightarrow_{\beta\eta}$.*

Corollary 12.3.8 Uniqueness of $\beta\eta$-Normal Forms. *Every untyped λ-term has at most one $\beta\eta$-normal form.*

Having thus disposed of uniqueness of $\beta\eta$-normal forms, we now turn to the problem of existence. We have already observed that the term $\Delta\Delta$ without β-normal form has no $\beta\eta$-normal form either, for the trivial reason that it has no η-redexes. We are now going to prove that the counterexample was forced on us, in the sense that the terms without $\beta\eta$-normal form are exactly the terms without β-normal form.

The main technical tool of the proof is the next result.

Proposition 12.3.9 Second Invertibility Property (Takahashi [1989]) *A parallel η-reduction followed by a parallel β-reduction can be inverted, in the sense of being replaced by a parallel β-reduction followed by a parallel η-reduction. Schematically:*

$$\frac{u \Longrightarrow_\eta \Longrightarrow_\beta v}{u \Longrightarrow_\beta \Longrightarrow_\eta v.}$$

Proof. For simplicity of notation, in the present proof we write

$$\lambda\vec{y}.\,(t)\vec{y}$$

as an abbreviation for

$$\lambda y_1.\,(\lambda y_2.\,\cdots(\lambda y_n.ty_n)\cdots y_2)y_1.$$

If $u \Longrightarrow_\eta t \Longrightarrow_\beta v$, we prove by induction on $t \Longrightarrow_\beta v$ that there is a term s such that $u \Longrightarrow_\beta s \Longrightarrow_\eta v$.

1. $u \Longrightarrow_\eta v \Longrightarrow_\beta v$
   Then $u \Longrightarrow_\beta u \Longrightarrow_\eta v$.

2. $u \Longrightarrow_\eta t_1t_2 \Longrightarrow_\beta v_1v_2$, with $t_1 \Longrightarrow_\beta v_1$ and $t_2 \Longrightarrow_\beta v_2$
   By definition of $\Longrightarrow_\eta$, $u = \lambda\vec{y}.\,(u_1u_2)\vec{y}$, with $u_1 \Longrightarrow_\eta t_1$ and $u_2 \Longrightarrow_\eta t_2$. Then

$$\dfrac{\dfrac{\dfrac{u_1 \Longrightarrow_\eta t_1 \Longrightarrow_\beta v_1}{u_1 \Longrightarrow_\beta s_1 \Longrightarrow_\eta v_1} \qquad \dfrac{u_2 \Longrightarrow_\eta t_2 \Longrightarrow_\beta v_2}{u_2 \Longrightarrow_\beta s_2 \Longrightarrow_\eta v_2}}{u_1u_2 \Longrightarrow_\beta s_1s_2 \Longrightarrow_\eta v_1v_2}}{\lambda\vec{y}.\,(u_1u_2)\vec{y} \Longrightarrow_\beta \lambda\vec{y}.\,(s_1s_2)\vec{y} \Longrightarrow_\eta v_1v_2}$$

   by the induction hypothesis and definition of $\Longrightarrow_\beta$ and $\Longrightarrow_\eta$.

3. $u \Longrightarrow_\eta \lambda x.\,t_1 \Longrightarrow_\beta \lambda x.\,v_1$, with $t_1 \Longrightarrow_\beta v_1$
   By definition of $\Longrightarrow_\eta$, $u = \lambda\vec{y}.\,(\lambda x.\,u_1)\vec{y}$, with $u_1 \Longrightarrow_\eta t_1$. Then

$$\dfrac{\dfrac{\dfrac{u_1 \Longrightarrow_\eta t_1 \Longrightarrow_\beta v_1}{u_1 \Longrightarrow_\beta s_1 \Longrightarrow_\eta v_1}}{\lambda x.\,u_1 \Longrightarrow_\beta \lambda x.\,s_1 \Longrightarrow_\eta \lambda x.\,v_1}}{\lambda\vec{y}.\,(\lambda x.\,u_1)\vec{y} \Longrightarrow_\beta \lambda\vec{y}.\,(\lambda x.\,s_1)\vec{y} \Longrightarrow_\eta \lambda x.\,v_1}$$

   by the induction hypothesis and definition of $\Longrightarrow_\beta$ and $\Longrightarrow_\eta$.

4. $u \Longrightarrow_\eta (\lambda x.\,t_1)t_2 \Longrightarrow_\beta v_1[x := v_2]$, with $t_1 \Longrightarrow_\beta v_1$ and $t_2 \Longrightarrow_\beta v_2$
   By definition of $\Longrightarrow_\eta$,

$$u = \lambda\vec{y}.\,((\lambda\vec{z}.\,(\lambda x.\,u_1)\vec{z})u_2)\vec{y},$$

with $u_1 \Longrightarrow_\eta t_1$ and $u_2 \Longrightarrow_\eta t_2$. By the induction hypothesis,

$$\frac{u_1 \Longrightarrow_\eta t_1 \Longrightarrow_\beta v_1}{u_1 \Longrightarrow_\beta s_1 \Longrightarrow_\eta v_1} \quad \text{and} \quad \frac{u_2 \Longrightarrow_\eta t_2 \Longrightarrow_\beta v_2}{u_2 \Longrightarrow_\beta s_2 \Longrightarrow_\eta v_2.}$$

We get the desired conclusion in two steps.

(a) We first notice that

$$\cfrac{\cfrac{u_1 \Longrightarrow_\beta s_1 \qquad u_2 \Longrightarrow_\beta s_2}{(\lambda \vec{z}.\,(\lambda x.\,u_1)\vec{z})u_2 \Longrightarrow_\beta s_1[x := s_2]}}{\lambda \vec{y}.\,((\lambda \vec{z}.\,(\lambda x.\,u_1)\vec{z})u_2)\vec{y} \Longrightarrow_\beta \lambda \vec{y}.\,(s_1[x := s_2])\vec{y}.}$$

The second step holds by definition of $\Longrightarrow_\beta$, while the first step is the only nontrivial part of the proof. We prove it for the case when $\vec{z}$ consists of z_1 and z_2, the general case being similar:

$$\cfrac{\cfrac{\cfrac{u_1 \Longrightarrow_\beta s_1 \qquad z_2 \Longrightarrow_\beta z_2}{(\lambda x.\,u_1)z_2 \Longrightarrow_\beta s_1[x := z_2]} \qquad z_1 \Longrightarrow_\beta z_1}{(\lambda z_2.\,(\lambda x.\,u_1)z_2)z_1 \Longrightarrow_\beta (s_1[x := z_2])[z_2 := z_1]} \qquad u_2 \Longrightarrow_\beta s_2}{(\lambda z_1.\,(\lambda z_2.\,(\lambda x.\,u_1)z_2)z_1)u_2 \Longrightarrow_\beta ((s_1[x := z_2])[z_2 := z_1])[z_1 := s_2],}$$

and the last term is simply $s_1[x := s_2]$, as needed.

(b) We then notice that

$$\cfrac{\cfrac{s_1 \Longrightarrow_\eta v_1 \qquad s_2 \Longrightarrow_\eta v_2}{s_1[x := s_2] \Longrightarrow_\eta v_1[x := s_2]}}{\lambda \vec{y}.\,(s_1[x := s_2])\vec{y} \Longrightarrow_\eta v_1[x := s_2].}$$

The second step holds by definition of $\Longrightarrow_\eta$, while the first step is a Substitution Lemma for $\Longrightarrow_\eta$ that can be easily proved (as in 8.5.3 for $\Longrightarrow_\beta$, and 8.6.13 for $\Longrightarrow_{\beta\eta}$). □

Notice that invertibility of parallel reductions holds only in the direction proved above, and not in the opposite one. For example,

$$\lambda x.\,(\lambda y.\,yx)z \Longrightarrow_\beta \lambda x.\,zx \Longrightarrow_\eta z,$$

but no η-reduction can be performed in the first term.

The next result, also called the **Postponement Theorem**, provides a normal form for $\beta\eta$-reductions in which all η-reductions follow all β-reductions. This is an analogue of 1.1.2, which provides a normal form for $\mathcal{N}$-proofs in which all →-introductions follow all →-eliminations.

Corollary 12.3.10 Normal Form of $\boldsymbol{\beta\eta}$-Reductions (Curry and Feys [1958]) *If $u \longrightarrow_{\beta\eta} v$, then there exists a term t such that $u \longrightarrow_\beta t \longrightarrow_\eta v$.*

Proof. A $\beta\eta$-reduction can be factored into steps of $1\beta\eta$-reductions, and hence into steps of parallel β-reductions and parallel η-reductions. By the invertibility proved in 12.3.9, this can be turned into a sequence of (parallel) β-reductions, followed by a sequence of (parallel) η-reductions. □

The next result shows that the strengthening of the Untyped Lambda Calculus provided by the η-rule has no effect on the existence of normal forms.

Theorem 12.3.11 (Curry, Hindley and Seldin [1972]) *A term has a $\beta\eta$-normal form if and only if it has a β-normal form.*

Proof. There are two directions to prove.

1. *if a term has a β-normal form, then it also has a $\beta\eta$-normal form*
   Suppose u is the β-normal form of the given term, and $u \Longrightarrow_\eta v$. We prove that v is still in β-normal form, so that performing η-reductions on a term in β-normal form does not produce any new β-redex. Since any η-reduction decreases the length of the given term, in finitely many steps we can reduce to a term which is *still* in β-normal form and *also* in η-normal form, i.e. it is in $\beta\eta$-normal form.

   The proof of the claim is immediate. If u is in β-normal form, then

   $$u = \lambda x_1 \cdots x_n . y u_1 \cdots u_m,$$

   with u_i in β-normal form. If $u \Longrightarrow_\eta v$, then either

   $$u = \lambda x_1 \cdots x_n . y v_1 \cdots v_m,$$

   with $u_i \Longrightarrow_\eta v_i$, or

   $$u = \lambda x_1 \cdots x_{n-1} . y v_1 \cdots v_{m-1},$$

   with $u_i \Longrightarrow_\eta v_i$. By the induction hypothesis, each v_i is in β-normal form because so is u_i. Then v is in β-normal form, too.

2. *if a term has a $\beta\eta$-normal form, then it also has a β-normal form*
   Suppose a term has a $\beta\eta$-normal form. By the Postponement Theorem, the latter can be reached by a sequence of β-reductions, followed by a sequence of η-reductions. The end result is obviously in β-normal form, being in $\beta\eta$-normal form. We prove that if $u \Longrightarrow_\eta v$ and v is in β-normal form, then u has a β-normal form. Then the term separating the β-reductions from the η-reductions still has a β-normal form, and hence so does any term reducible to it.

Suppose $u \Longrightarrow_\eta v$ and v is in β-normal form. Then

$$v = \lambda x_1 \cdots x_n .\, y v_1 \cdots v_m,$$

with v_i in β-normal form. For simplicity of notation we restrict to the case

$$v = \lambda x.\, y v_1,$$

with v_1 in β-normal form, the general case being similar. If $u \Longrightarrow_\eta v$, then

$$u = \lambda \vec{z}_3.\, (\lambda x.\, (\lambda \vec{z}_2.\, ((\lambda \vec{z}_1.\, (y)\vec{z}_1)u_1)\vec{z}_2))\vec{z}_3,$$

with $u_1 \Longrightarrow_\eta v_1$. By the induction hypothesis, u_1 has a β-normal form. As in the proof of case 4.(a) of 12.3.9, by successive substitutions we can prove first

$$\lambda \vec{z}_2.\, ((\lambda \vec{z}_1.\, (y)\vec{z}_1)u_1)\vec{z}_2 \Longrightarrow_\beta \lambda \vec{z}_2.\, (y u_1)\vec{z}_2,$$

and then

$$u \Longrightarrow_\beta \lambda x.\, (\lambda \vec{z}_2.\, (y u_1)\vec{z}_2).$$

Since u_1 has a β-normal form, so does u. □

Normalizing reduction strategies

12.3.10 and 12.3.11 immediately allow us to exhibit the following $\beta\eta$-normalizing strategy:

1. reduce the leftmost β-redex, if there is one
2. reduce the leftmost η-redex, otherwise.

If the $\beta\eta$-normal form exists, then the first part produces the β-normal form and the second part produces its η-normal form, i.e. the $\beta\eta$-normal form.

This is not a literal leftmost reduction strategy yet, since the leftmost β-redex may not be the leftmost $\beta\eta$-redex, if there is an η-redex to its left. To make the notion precise, we extend the definition 12.2.9 of $\longrightarrow_{l\beta}$ as follows.

Definition 12.3.12 *A one-step $\beta\eta$-reduction $u \longrightarrow_{1\beta\eta} v$ is called a* **leftmost $\beta\eta$-reduction** *($u \longrightarrow_{1l\beta\eta} v$) if it reduces the leftmost β-redex or η-redex that can be reduced.*

We define $\longrightarrow_{l\beta\eta}$ as in 12.1.5, i.e. as the reflexive and transitive closure of $\longrightarrow_{1l\beta\eta}$.

Theorem 12.3.13 Leftmost $\beta\eta$-Reduction Strategy (Klop [1980]) *If $u \longrightarrow_{\beta\eta} v$ and v is in $\beta\eta$-normal form, the $u \longrightarrow_{l\beta\eta} v$.*

Proof. By 12.3.10, there is a term t such that

$$u \longrightarrow_{\beta} t \longrightarrow_{\eta} v.$$

The β-reduction can turned into a leftmost β-reduction by 12.2.13, since t is in β-normal form by the proof of 12.3.11.2. The η-reduction can be turned into a leftmost $\beta\eta$-reduction by always reducing the leftmost η-redex, since all its terms are in β-normal form by the proof of 12.3.11.1. Thus

$$u \longrightarrow_{l\beta} t \longrightarrow_{l\beta\eta} v.$$

We prove by induction on the total length that

$$\longrightarrow_{l\beta}\longrightarrow_{l\beta\eta} = \longrightarrow_{l\beta\eta} .$$

Consider the leftmost β-reduction. If it is already a leftmost $\beta\eta$-reduction, then

$$\longrightarrow_{l\beta}\longrightarrow_{l\beta\eta} = \longrightarrow_{1\beta\eta} (\longrightarrow_{l\beta}\longrightarrow_{l\beta\eta}) \ = \longrightarrow_{l\beta\eta}$$

by the induction hypothesis applied to the reduction in parenthesis.

Otherwise, let $\lambda x.\, u_1 x$ be the leftmost $\beta\eta$-redex. There are three cases:

1. The leftmost β-redex is $u_1 x$, i.e. $u_1 = \lambda y.\, u_2$. Since

   $$\ldots \lambda x.\, (\lambda y.\, u_2)x \ldots \longrightarrow_{l\beta} \ldots \lambda x.\, u_2[y := x] \ldots = \ldots \lambda y.\, u_2 \ldots = \ldots u_1 \ldots,$$

   we can obtain the same result by performing first the leftmost $\beta\eta$-reduction:

   $$\ldots \lambda x.\, u_1 x \ldots \longrightarrow_{1l\beta\eta} \ldots u_1 \ldots$$

   This decreases the length of the initial $\longrightarrow_{l\beta}$.

2. The leftmost β-redex is inside u_1, and $u_1 \longrightarrow_{l\beta} u_1'$. Since

   $$\ldots \lambda x.\, u_1 x \ldots \longrightarrow_{l\beta} \ldots \lambda x.\, u_1' x \ldots \longrightarrow_{1\eta} \ldots u_1' \ldots,$$

   we can obtain the same result by performing first the leftmost $\beta\eta$-reduction:

   $$\ldots \lambda x.\, u_1 x \ldots \longrightarrow_{1l\beta\eta} \ldots u_1 \ldots \longrightarrow_{l\beta} \ldots u_1' \ldots$$

   This decreases the length of the final $\longrightarrow_{l\beta\eta}$.

3. The leftmost β-redex s is after $\lambda x.\, u_1 x$, and $s \longrightarrow_{l\beta} s'$. Since

   $$\ldots \lambda x.\, u_1 x \ldots s \ldots \longrightarrow_{l\beta} \ldots \lambda x.\, u_1 x \ldots s' \ldots,$$

   we can obtain the same result by performing first the leftmost $\beta\eta$-reduction:

   $$\ldots \lambda x.\, u_1 x \ldots s \ldots \longrightarrow_{1l\beta\eta} \ldots u_1 \ldots s \ldots \longrightarrow_{l\beta} \ldots u_1 \ldots s' \ldots$$

   This decreases the length of the final $\longrightarrow_{l\beta\eta}$. □

æ

Chapter 13

Semantics

We now turn to semantical interpretations of the Untyped Lambda Calculus, by following the blueprint of Chapter 10.

The type structure allowed us to obtain models of the Typed Lambda Calculus quite easily. When types are missing the definition of a model is simpler, although the construction of a model is more complicated.

Definition 13.0.14 *A* **model** *of the Untyped Lambda Calculus is a structure*

$$\mathcal{M} = \langle M, [\![\ \ ]\!]^{\mathcal{M}} \rangle$$

with the following properties, where ρ is an environment for M:

1. *if t is a term, then $[\![t]\!]^{\mathcal{M}}_{\rho} \in M$.*
2. *$[\![\ \ ]\!]^{\mathcal{M}}$ respects β-equality, i.e.*

   $$t_1 =_{\beta} t_2 \;\Rightarrow\; (\forall \rho)([\![t_1]\!]^{\mathcal{M}}_{\rho} = [\![t_2]\!]^{\mathcal{M}}_{\rho}),$$

   where the equality on the right indicates identity of objects in M.

An **extensional model** *is a model that respects $\beta\eta$ equality, i.e.*

$$t_1 =_{\beta\eta} t_2 \;\Rightarrow\; (\forall \rho)([\![t_1]\!]^{\mathcal{M}}_{\rho} = [\![t_2]\!]^{\mathcal{M}}_{\rho}).$$

To improve readability we can omit either of the indeces ρ or $\mathcal{M}$ in $[\![\ \ ]\!]^{\mathcal{M}}_{\rho}$, when no confusion arises.

Moreover, we will also write

$$\mathcal{M} \models t_1 = t_2 \qquad \text{for} \qquad (\forall \rho)([\![t_1]\!]^{\mathcal{M}}_{\rho} = [\![t_2]\!]^{\mathcal{M}}_{\rho}),$$

so that the soundness conditions in the definition of a model can be stated more succinctly as

$$t_1 =_\beta t_2 \;\Rightarrow\; \mathcal{M} \models t_1 = t_2$$

or

$$t_1 =_{\beta\eta} t_2 \;\Rightarrow\; \mathcal{M} \models t_1 = t_2.$$

13.1 Term Models

Since normal forms do not always exist, there is no analogue for the Untyped Lambda Calculus of the first term model $\mathcal{T}_1$ of Chapter 10. The following provides an analogue of the second term model $\mathcal{T}_2$.

Definition 13.1.1 *The* **term model** $\mathcal{T}$ *is defined as follows:*

1. *The underlying structure consists of:*

$$T = \{\textit{equivalence classes of terms w.r.t. } =_\beta\}$$

2. *Given an environment ρ on T, i.e. a function assigning to every variable an equivalence class, let ρ^* be a choice function for ρ, i.e. a function that associates to every variable x a term in the equivalence class $\rho(x)$. Then*

$$[\![t]\!]^{\mathcal{T}}_\rho = \textit{the equivalence class of } t[\vec{x} := \rho^*(\vec{x})],$$

where $t[\vec{x} := \rho^*(\vec{x})]$ indicates the result of the simultaneous substitution of the term $\rho^*(x)$ for any free variable x of t.

Proposition 13.1.2 *The structure $\mathcal{T}$ is a model of the Untyped Lambda Calculus. Actually,*

$$t_1 =_\beta t_2 \;\Leftrightarrow\; \mathcal{T} \models t_1 = t_2.$$

Proof. If $t_1 =_\beta t_2$, then

$$t_1[\vec{x} := \rho^*(\vec{x})] =_\beta t_2[\vec{x} := \rho^*(\vec{x})],$$

because $=_\beta$ is invariant under simultaneous substitutions. But β-equal terms are in the same equivalence class, and thus

$$[\![t_1]\!]_\rho = [\![t_2]\!]_\rho.$$

Conversely, suppose $t_1 \neq_\beta t_2$. Then they are in different equivalence classes. If ρ is the environment associating every variable to its equivalence class, we can choose as ρ^* the identity function. Then

$$t_1[\vec{x} := \rho^*(\vec{x})] = t_1 \neq_\beta t_2 = t_2[\vec{x} := \rho^*(\vec{x})],$$

and hence

$$[\![t_1]\!]_\rho \neq [\![t_2]\!]_\rho$$

because $t_1[\vec{x} := \rho^*(\vec{x})]$ and $t_2[\vec{x} := \rho^*(\vec{x})]$ are in different equivalence classes. □

An extensional model can be defined similarly, by considering $\beta\eta$-equality.

As usual, the term models provide semantical interpretations of terms too close to the original syntactical presentation. In other words, they are well-behaved but not very insightful. This is the reason to continue, in the next sections, the search for other more informative models.

13.2 Functional Models

We now turn to the consideration of models in which terms $\lambda x.\, u$ are interpreted as functions in the usual mathematical sense. We have already considered such models for the Typed Lambda Calculus, but here there is an additional complication, due to the fact that the same term can behave as a function or as an argument, according to the situation. Since the domain M of the model is fixed, we need a way of treating elements of M as functions on M, and conversely.

In practice, we define the interpretation function $[\![\ ]\!]^{\mathcal{M}}$ in the following canonical way, which reduces it to the *definition* of two functions I and J, identifying elements of M with functions on M and conversely, and to the *verification* of closure under informal abstraction.

Definition 13.2.1 Canonical Interpretation. *Given M, together with a total function*

$$I : M \to M^M$$

and a partial function

$$J : M^M \to M,$$

and an environment ρ on M, we define $[\![\ ]\!]_\rho$ by induction on terms, as follows:

$$[\![t]\!]_\rho = \begin{cases} \rho(x) & \text{if } t = x \\ I([\![u]\!]_\rho)([\![v]\!]_\rho) & \text{if } t = uv \\ J(\Lambda X.\, [\![u]\!]_{\rho[x:=X]}) & \text{if } t = \lambda x.\, u, \end{cases}$$

where $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ denotes the function

$$a \in M \longmapsto [\![u]\!]_{\rho[x:=a]} \in M.$$

We should check, inductively, that $[\![t]\!]_\rho$ is a member of M. Or, equivalently, that $[\![t]\!]_\rho$ is defined for every term t. In the first clause, it is so by definition of environment. In the second clause, it is so by the induction hypothesis, since then

$I(\llbracket u \rrbracket_\rho)$ is a function on M. In the last case, by the induction hypothesis, we obtain a function from M to M, but not necessarily a function in the domain of J, without further hypotheses on J. For specific structures, this will have to be *verified*.

Not requiring *totality for* J corresponds to not requiring full functional models for the Typed Lambda Calculus. But, while there the condition was optional (see 10.3.4), here it will turn out to be necessary (see 13.2.4).

The next results is central to the later development, and its idea is simple. Since the β-rule eliminates the application of a λ-abstraction, in a model $I \circ J$ should cancel. Similarly, since the η-rule eliminates the λ-abstraction of an application, in an extensional model $J \circ I$ should cancel.

Proposition 13.2.2 Sufficient Conditions for Models. *The structure*

$$\mathcal{M} = \langle M, I, J, \llbracket\ \ \rrbracket^{\mathcal{M}} \rangle,$$

where $\llbracket\ \ \rrbracket^{\mathcal{M}}$ *is the canonical interpretation, is:*

1. *a model of the Untyped Lambda Calculus if* $I \circ J$ *is the identity (on the domain of* J*)*

2. *an extensional model of the Untyped Lambda Calculus if, moreover,* $J \circ I$ *is the identity (on* M*).*

Proof. The proof is by induction on $\longrightarrow_{\beta\eta}$. We only consider the two crucial cases of the reduction rules, the other ones being as in 10.3.4.

1. $(\lambda x.\, u)v \longrightarrow_{1\beta\eta} u[x := v]$
   Then

$$\begin{aligned}
\llbracket (\lambda x.\, u)v \rrbracket_\rho &= I(\llbracket \lambda x.\, u \rrbracket_\rho)(\llbracket v \rrbracket_\rho) \\
&= I(J(\Lambda X.\, \llbracket u \rrbracket_{\rho[x:=X]}))(\llbracket v \rrbracket_\rho) \\
&= (\Lambda X.\, \llbracket u \rrbracket_{\rho[x:=X]})(\llbracket v \rrbracket_\rho) \\
&= \llbracket u \rrbracket_{\rho[x:=\llbracket v \rrbracket_\rho]} \\
&= \llbracket u[x := v] \rrbracket_\rho.
\end{aligned}$$

   The first two equalities hold by definition of $\llbracket\ \ \rrbracket$. The third holds because $I \circ J$ is the identity. The fourth holds by definition of Λ as a mathematical function, which is computed by instantiating the variable to the argument. The last equality holds by the fact, easily proved by induction on u, that substitution in terms and in environments commute.

2. $\lambda x.\, tx \longrightarrow_{1\beta\eta} t$
   Then

$$\llbracket \lambda x.\, tx \rrbracket_\rho \;=\; J(\Lambda X.\, \llbracket tx \rrbracket_{\rho[x:=X]})$$

$$\begin{aligned}
&= J(\Lambda X.\, I([\![t]\!]_{\rho[x:=X]})([\![x]\!]_{\rho[x:=X]}))\\
&= J(\Lambda X.\, I([\![t]\!]_{\rho[x:=X]})(X))\\
&= J(I([\![t]\!]_{\rho}))\\
&= [\![t]\!]_{\rho}.
\end{aligned}$$

The first three equalities hold by definition of $[\![\ \ ]\!]$. The fourth holds by definition of Λ. The last holds because $J \circ I$ is the identity. □

Notice that *if $I \circ J$ is the identity, then J is one-one and I is onto the domain of J.* One-oneness follows from the fact that

$$J(x_1) = J(x_2) \;\Rightarrow\; I(J(x_1)) = I(J(x_2)) \;\Rightarrow\; x_1 = x_2.$$

Ontoness follows from the fact that $J(x)$ is a counterimage of x in the domain of J under I, since $I(J(x)) = x$.

A trivial modification of the proof of 13.2.2 shows that Condition 1, i.e. that $I \circ J$ be the identity, implies the following *weak form of extensionality*:[1]

$$t_1 =_{\beta\eta} t_2 \;\Rightarrow\; (\forall\rho)(I([\![t_1]\!]_{\rho}) = I([\![t_2]\!]_{\rho})).$$

The full version of extensionality, i.e.

$$t_1 =_{\beta\eta} t_2 \;\Rightarrow\; (\forall\rho)([\![t_1]\!]_{\rho} = [\![t_2]\!]_{\rho}),$$

follows from the weak form if I is one-one, because then

$$I([\![t_1]\!]_{\rho}) = I([\![t_2]\!]_{\rho}) \;\Rightarrow\; [\![t_1]\!]_{\rho} = [\![t_2]\!]_{\rho}.$$

This is precisely what Condition 2 ensures, since *if $I \circ J$ is the identity, then I is one-one if and only if $J \circ I$ is the identity.* One direction follows as above. For the other direction, suppose I is one-one. Given $x \in M$, let $y = J(I(x))$. Then

$$I(y) = I(J(I(x))) = I(x) \;\Rightarrow\; y = x.$$

Obviously, to get an extensional model we would only need I to be one-one on the elements of M that are interpretations of terms, while Condition 2 requires I to be one-one on all elements of M. Although stronger, this is obviously easier to deal with, because it avoids the problem of knowing which elements of M are interpretations of terms.

Before turning to a search for models, we notice that simple extremal solutions would not work here, as they did instead for the Typed Lambda Calculus, where we obtained full functional models over any set A, in particular nontrivial models whose levels are all finite.

[1] The notion of a model we adopted is called a λ-*algebra* in the literature. A weakly extensional model in the sense above is called a λ-*model*.

Proposition 13.2.3 Impossibility of Nontrivial Finite Functional Models. *There is no finite functional model of the Untyped Lambda Calculus with at least two elements.*

Proof. Consider any finite functional model. Since there is only a fixed number of interpretations, but for any n there are n terms of the form $\lambda x_1 \cdots x_n .\, x_i$, with $1 \leq i \leq n$, there must exist n and $1 \leq i < j \leq n$ such that the interpretations of $\lambda x_1 \cdots x_n .\, x_i$ and $\lambda x_1 \cdots x_n .\, x_j$ are the same.

Since the model uses the canonical interpretation of terms,

$$[\![\lambda x_1 \cdots x_n .\, x_i]\!] = \Lambda X_1 \cdots \Lambda X_n .\, X_i \quad \text{and} \quad [\![\lambda x_1 \cdots x_n .\, x_j]\!] = \Lambda X_1 \cdots \Lambda X_n .\, X_j .$$

Given any two elements a and b of the model, it is enough to choose any string $(c_1, \ldots, c_n)$ of elements such that $c_i = a$ and $c_j = b$ to have

$$a = [\![\lambda x_1 \cdots x_n .\, x_i]\!](c_1, \ldots, c_n) = [\![\lambda x_1 \cdots x_n .\, x_j]\!](c_1, \ldots, c_n) = b.$$

Then the model is trivial, because all the elements coincide. □

Thus, while a nontrivial functional model of the Typed Lambda Calculus must have at least two distinct interpretations of numerals, but it can have exactly two (see 11.3.3), *a nontrivial functional model of the Untyped Lambda Calculus must give distinct interpretations to all numerals*.

Proposition 13.2.4 Impossibility of Nontrivial Full Functional Models (Cantor [1874]) *There is no set M with at least two elements such that $I \circ J$ is the identity on M^M, where*

$$I : M \to M^M \qquad \textit{and} \qquad J : M^M \to M$$

are both total functions.

Proof. Suppose M has two distinct elements a and b, and let

$$d(x) = \begin{cases} a & \text{if } I(x)(x) \neq a \\ b & \text{otherwise.} \end{cases}$$

be a diagonal function. If J is total, $J(d)$ is defined and

$$d(J(d)) = \begin{cases} a & \text{if } I(J(d))(J(d)) \neq a \\ b & \text{otherwise.} \end{cases}$$

But $I \circ J$ is the identity, and thus $I(J(d)) = d$. Then

$$d(J(d)) = \begin{cases} a & \text{if } d(J(d)) \neq a \\ b & \text{otherwise,} \end{cases}$$

contradiction. □

Since there is no hope of finding models with a total J, the whole point of the work to come will be to isolate appropriate subsets of M^M as domains for J. By so doing, the verification that $[\![u]\!]_\rho \in M$ will become nontrivial, since the function $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ used in the last clause of the canonical definition of $[\![\ ]\!]_\rho$ will have to be in the domain of J.

13.3 The Graph Model

13.2.4 shows that to obtain a functional model of the Untyped Lambda Calculus we have to impose some restrictions on the kind of functions we use as interpretations of terms. When dealing with the Typed Lambda Calculus we discovered a restriction that worked, i.e. *continuity*. We thus look for a c.c.p.o. D and total[2] continuous functions

$$I : D \to [D \to D] \qquad \text{and} \qquad J : [D \to D] \to D$$

such that $I \circ J = id_{[D \to D]}$. Continuity of I and J will allow us to extend the proof of 10.4.3, and prove that the canonical interpretation 13.2.1 is well-defined, i.e. that $[\![t]\!]_\rho \in D$ for any environment ρ on D. The condition on $I \circ J$ will instead allow us to apply 13.2.2, and claim that D is a (nonextensional) model.

The intuition behind the construction of the graph model is the following. A continuous function on the reals can easily be coded by a single real, because such a function is completely determined by its values on the rationals, and countably many reals can be coded by a single real. This provides a suggestion for the definition of I.

Conversely, any real can be thought of as the code of countably many reals, and hence of the values of a function on the rationals. Unfortunately, assigning arbitrary real values to the rationals does not in general determine a continuous function on the reals, because their limits may not exist. We thus have an obstacle for the definition of J.

We then step from the reals to $\mathcal{P}(\omega)$, the set of all sets of natural numbers. A continuous function on $\mathcal{P}(\omega)$ can easily be coded by a single set of natural numbers, because such a function is completely determined by its values on the finite sets, and countably many sets of natural numbers can be coded by a single set. This retains the suggestion for the definition of I.

Conversely, any set of natural numbers can be thought of as the code of countably many sets, and hence of the values of a function on the finite sets. Fortunately, assigning arbitrary values to the finite sets does determine a continuous function

[2] We argued in 13.2.4 that J cannot be total on D^D. Thus the intended domain $[D \to D]$ of J, to which the word ‘total’ refers to here, will be a proper subset of D^D.

on $\mathcal{P}(\omega)$, because their limits are simply their unions. This provides a suggestion for the definition of J.

It now remains to put the plan into practice.

Continuous functions on sets

Recall that $\mathcal{P}(\omega)$ is a c.c.p.o. (actually, a c.p.o.), with set-theoretical inclusion as the ordering, and set-theoretical union as the l.u.b. operation.

We already know from Exercise 6.3.9.c that the continuous functions on $\mathcal{P}(\omega)$ are exactly the functions determined by their behavior on finite sets. Since this is the crucial property that allows us to consider $\mathcal{P}(\omega)$ as a model of the Untyped Lambda Calculus, we reprove it directly here.

For simplicity of notations, *in the present section we use the letter u as a variable on finite sets of natural numbers*.

Proposition 13.3.1 (Uspenskii [1955], Nerode [1957]) *A function $f : \mathcal{P}(\omega) \to \mathcal{P}(\omega)$ is continuous if and only if, for every $A \subseteq \omega$,*

$$f(A) = \bigcup_{u \subseteq A} f(u).$$

Proof. If f is continuous and A is any subset of ω, then

$$u \subseteq A \;\Rightarrow\; f(u) \subseteq f(A)$$

by monotonicity. Thus

$$\bigcup_{u \subseteq A} f(u) \subseteq f(A).$$

Conversely, let $A_n = A \cap \{0, \dots, n\}$. Then $A = \bigcup_{n \in \omega} A_n$, and by continuity

$$f(A) = f(\bigcup_{n \in \omega} A_n) = \bigcup_{n \in \omega} f(A_n).$$

Since each A_n is a finite subset of A,

$$f(A) \subseteq \bigcup_{u \subseteq A} f(u).$$

Putting together the two inequalities just proved, we obtain

$$f(A) = \bigcup_{u \subseteq A} f(u).$$

Conversely, suppose

$$f(A) = \bigcup_{u \subseteq A} f(u)$$

for any A. First, f is monotone because if $A \subseteq B$, then every u contained in A is also contained in B, and thus $f(A) \subseteq f(B)$ because the latter is a bigger union.

Let now $A_0 \subseteq A_1 \subseteq \cdots$ be a given chain of arbitrary (not necessarily finite) sets. Then

$$\begin{array}{rcl} f(\bigcup_{n\in\omega} A_n) & = & \bigcup\{f(u) : u \subseteq \bigcup_{n\in\omega} A_n\} \\ & = & \bigcup_{n\in\omega} \bigcup\{f(u) : u \subseteq A_n\} \\ & = & \bigcup_{n\in\omega} f(A_n) \end{array}$$

by the hypothesis on f, the fact that a finite set is contained in $\bigcup_{n\in\omega} A_n$ if and only if it is contained in some A_n (because the A_n form a chain), and the hypothesis on f again. □

Coding of finite sets

We now know that to code a continuous function on $\mathcal{P}(\omega)$ we only need to code its behavior on the finite sets. This requires the following steps:

1. *coding a finite set of numbers by a single number*
   We identify numbers and finite sets as follows:

   $$u_m = \{x_1, \dots, x_n\} \Leftrightarrow m = 2^{x_1} + \cdots + 2^{x_n}.$$

   By convention, $u_0 = \emptyset$.

   Note that every number different from 0 can be uniquely decomposed in the form $2^{x_1} + \cdots + 2^{x_n}$ with all x_i distinct, and thus the coding is actually one-one and onto.

   Intuitively, a number m codes a finite set as follows. If m in written in binary notation, then u_m consists of the elements corresponding to the positions of 1's from right to left. For example, $77 = 1001101$ codes the set $u_{77} = \{0, 2, 3, 6\}$.

2. *coding an ordered pair of numbers by a single number*
   We associate to each pair (n, m) of natural numbers the single number

   $$\langle n, m\rangle = 2^n(2m + 1) - 1.$$

Note that every number different from 0 can be uniquely written as the product of an even (2^n) and an odd ($2m+1$) number, and thus the coding is actually one-one and onto.

3. *coding a sequence of sets by a single set*
We associate to every sequence A_n of sets a single set $\bigoplus_{n\in\omega} A_n$ defined as follows:
$$\bigoplus_{n\in\omega} A_n = \{\langle x, n\rangle : x \in A_n\}.$$
Note that every set A can be uniquely decomposed in the sequence
$$(A)_n = \{x : \langle x, n\rangle \in A\},$$
and thus the coding is actually one-one and onto. Moreover,
$$A = \bigoplus_{n\in\omega}(A)_n \qquad \text{and} \qquad A_n = (\bigoplus_{m\in\omega} A_m)_n.$$

A nonextensional model

We are now ready for our first model.

Theorem 13.3.2 The Graph Model (Plotkin [1972], Scott [1975]) $\mathcal{P}(\omega)$ *induces a (nonextensional) canonical model of the Untyped Lambda Calculus.*

Proof. We define the functions
$$I : \mathcal{P}(\omega) \to [\mathcal{P}(\omega) \to \mathcal{P}(\omega)] \qquad \text{and} \qquad J : [\mathcal{P}(\omega) \to \mathcal{P}(\omega)] \to \mathcal{P}(\omega)$$
in the natural way, as follows.

- Given a continuous function $f \in [\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$, we identify it with the set coding the sequence of its values for all finite arguments:
$$J(f) = \bigoplus_{n\in\omega} f(u_n).$$

- Given a set $A \in \mathcal{P}(\omega)$, we consider its n-th component as the value given to the n-th finite set by the function coded by A:[3]
$$I(A) = \Lambda X. \bigcup_{u_n \subseteq X} (A)_n.$$

[3] Despite the intuition just described, it is *not* the case that $I(A)(u_n) = (A)_n$. Rather, $I(A)(u_n) = \bigcup\{(A)_m : u_m \subseteq u_n\}$. This is needed for the continuity of $I(A)$, proved below.

We have to check that I is well-defined, i.e. that $I(A)$ *is a continuous function on* $\mathcal{P}(\omega)$*, for any* A. Let

$$X_0 \subseteq X_1 \subseteq \cdots$$

be a chain in $\mathcal{P}(\omega)$. Then

$$\begin{array}{rcl} I(A)(\bigcup_{m\in\omega} X_m) & = & \bigcup\{(A)_n : u_n \subseteq \bigcup_{m\in\omega} X_m\} \\ & = & \bigcup_{m\in\omega} \bigcup\{(A)_n : u_n \subseteq X_m\} \\ & = & \bigcup_{m\in\omega} I(A)(X_m) \end{array}$$

by definition of $I(A)$, the fact that a finite set is contained in $\bigcup_{m\in\omega} X_m$ if and only if it is contained in some X_m (because the X_m form a chain), and definition of $I(A)$ again.

Consider now the structure

$$\langle \mathcal{P}(\omega), I, J, [\![\;\;]\!]^{\mathcal{P}(\omega)} \rangle,$$

where $[\![\;\;]\!]^{\mathcal{P}(\omega)}$ is the canonical interpretation defined in 13.2.1. Recall that, given any environment ρ on $\mathcal{P}(\omega)$,

$$[\![t]\!]_\rho = \left\{ \begin{array}{ll} \rho(x) & \text{if } t = x \\ I([\![u]\!]_\rho)([\![v]\!]_\rho) & \text{if } t = uv \\ J(\Lambda X.\, [\![u]\!]_{\rho[x:=X]}) & \text{if } t = \lambda x.\, u, \end{array} \right.$$

where $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ denotes the function

$$A \in \mathcal{P}(\omega) \longmapsto [\![u]\!]_{\rho[x:=A]} \in \mathcal{P}(\omega).$$

By 13.2.2, to prove that $\mathcal{P}(\omega)$ is a model it is enough to check that $I \circ J$ is the identity on $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$, and that $[\![\;\;]\!]_\rho$ is well-defined, i.e. that $[\![t]\!]_\rho \in \mathcal{P}(\omega)$ for any t and ρ. We prove the former below. For the latter, the only differences with the proof of 10.4.3 are the absence of types, and the presence of I and J in the definition of $[\![\;\;]\!]_\rho$. The first presents no problem, and the second is taken care of by the fact, proved below, that I and J are continuous, which makes the proof of the fact that $[\![u]\!]_{\rho[x:=X]}$ is a continuous function of X go through.

The following facts conclude the proof.

1. *$I \circ J$ is the identity*
   Given $f \in [\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ and $X \in \mathcal{P}(\omega)$,

$$I(J(f))(X) \quad = \quad \bigcup_{u_n \subseteq X} (J(f))_n$$

$$\begin{aligned}
&= \bigcup_{u_n \subseteq X} (\bigoplus_{m\in\omega} f(u_m))_n \\
&= \bigcup_{u_n \subseteq X} f(u_n) \\
&= f(X)
\end{aligned}$$

by the definitions of I and J, properties of the coding, and 13.3.1 (which holds because f is continuous).

2. *I is continuous*
If

$$A_0 \subseteq A_1 \subseteq \cdots$$

is a chain in $\mathcal{P}(\omega)$, then

$$\begin{aligned}
I(\bigcup_{m\in\omega} A_m)(X) &= \bigcup_{u_n\subseteq X} (\bigcup_{m\in\omega} A_m)_n \\
&= \bigcup_{u_n\subseteq X} \bigcup_{m\in\omega} (A_m)_n \\
&= \bigcup_{m\in\omega} \bigcup_{u_n\subseteq X} (A_m)_n \\
&= \bigcup_{m\in\omega} I(A_m)(X) \\
&= (\bigsqcup_{m\in\omega}^{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]} I(A_m))(X)
\end{aligned}$$

by the definition of I, properties of union, and the definition of $\bigsqcup_{m\in\omega}^{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]}$.

3. *J is continuous*
If

$$f_0 \sqsubseteq_{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]} f_1 \sqsubseteq_{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]} \cdots$$

is a chain in $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$, then

$$\begin{aligned}
J(\bigsqcup_{m\in\omega}^{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]} f_m) &= \bigoplus_{n\in\omega}(\bigsqcup_{m\in\omega}^{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]} f_m)(u_n) \\
&= \bigoplus_{n\in\omega}\bigcup_{m\in\omega} f_m(u_n) \\
&= \{\langle x,n\rangle : x \in \bigcup_{m\in\omega} f_m(u_n)\}
\end{aligned}$$

$$\begin{aligned} &= \bigcup_{m\in\omega} \{\langle x, n\rangle : x \in f_m(u_n)\} \\ &= \bigcup_{m\in\omega} \bigoplus_{n\in\omega} f_m(u_n) \\ &= \bigcup_{n\in\omega} J(f_m) \end{aligned}$$

by the definitions of J, $\bigsqcup_{m\in\omega}^{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]}$ and $\bigoplus$, properties of union, and the definitions of $\bigoplus$ and J. □

The fact that $I \circ J$ is the identity simply says that the function associated with the set coding f is f itself, and holds by the proof above.

The fact that $J \circ I$ is the identity says that the set associated with the function coded by a set X is X itself, and fails because the same function can be coded by different sets. For example, the identity function is coded by both the following sets:

$$A = \bigoplus_{n\in\omega} u_n \qquad \text{and} \qquad B = \bigoplus_{m\in\omega} u_{2^m},$$

because a set X is both the union of all finite sets contained in it, and of all singletons contained in it (recall that $u_{2^m} = \{m\}$). But

$$\begin{aligned} J(I(B)) &= J(\Lambda X. \bigcup_{u_n\subseteq X} (B)_n) \\ &= J(\Lambda X. \bigcup_{m\in X} \{m\}) \\ &= J(\Lambda X. X) \\ &= \bigoplus_{n\in\omega} u_n \\ &= A \neq B. \end{aligned}$$

In general, only $J(I(A)) \supseteq A$ holds, for any A.

Thus $\mathcal{P}(\omega)$ satisfies Condition 1 of 13.2.2 for a model, but not Condition 2 for an extensional model. Actually, $\mathcal{P}(\omega)$ *is not an extensional model*, because $\lambda x. x$ and $\lambda yx.xy$ are extensionally equal but have different interpretations. More precisely, $[\![\lambda x. x]\!] = \Lambda X. X$ and $[\![\lambda yx.xy]\!] = \Lambda Y. \Lambda X. I(X)(Y)$.

At this point an *extensional model* could be obtained quite elegantly, using the method of retracts of Section 13.5. The price for elegance would be a lack of motivation, that comes from the work done in Section 13.4 for D_∞. However, the impatient reader can skip Section 13.4 and turn immediately to Section 13.5, that can be read independently.

Exercises 13.3.3 A coding-free version of $\mathcal{P}(\omega)$ (Plotkin [1972], Engeler [1981]) When dealing with $\mathcal{P}(\omega)$ we are forced to talk indirectly about the pair (x, u_n), consisting of a number x and a finite set u_n, through the number $\langle x, n\rangle$, obtained by a double coding (first of finite sets by numbers, and then of pairs of numbers by a single number). We can obtain a slight simplification of the Graph Model by replacing ω by any set Ω closed under pairs consisting of an element x and a finite set u, which allows us to talk directly about the pair (x, u).

a) *For any set B, there is a set $\Omega_B \supseteq B$ closed under pairs consisting of an element and a finite set.* (Hint: let $B_0 = B$, $B_{n+1} = B_n \cup \{(b, u) : b \in B_n \wedge u \subseteq B_n\}$ and $\Omega_B = \bigcup_{n\in\omega} B_n$.)

b) *For any set B, $\mathcal{P}(\Omega_B)$ is a c.c.p.o.*

c) *For any set B, $\mathcal{P}(\Omega_B)$ is a (nonextensional) canonical model of the Untyped Lambda Calculus.* (Hint: define the functions

$$I : \mathcal{P}(\Omega_B) \to [\mathcal{P}(\Omega_B) \to \mathcal{P}(\Omega_B)] \quad \text{and} \quad J : [\mathcal{P}(\Omega_B) \to \mathcal{P}(\Omega_B)] \to \mathcal{P}(\Omega_B)$$

in the natural way, as follows:

$$I(A) = \Lambda X. \bigcup_{u\subseteq X} \{x : (x, u) \in A\} \quad \text{and} \quad J(f) = \{(x, u) : x \in f(u)\}.$$

Then proceed as in 13.3.2, by proving in particular that $I \circ J$ is the identity, and that I and J are continuous.)

The Effective Graph Model ⋆

In the Graph Model $\mathcal{P}(\omega)$ there is a disparity between the countable number of λ-terms to be interpreted, and the uncountable number of elements available for the interpretation. However, it is easy to carve inside $\mathcal{P}(\omega)$ a countable subset that already provides a model. The idea is to restrict attention to the set $\mathcal{E}$ of the *recursively enumerable (r.e.) sets*. The computability notions needed for the present subsection can be found in Odifreddi [1989], to which we refer for background and notation.

The next definition isolates the appropriate continuous functions needed for the model.

Definition 13.3.4 (Friedberg and Rogers [1959]) *A continuous function f on $\mathcal{P}(\omega)$ is* **effective** *if its graph $J(f)$ is an r.e. set.*[4]

Proposition 13.3.5 *The effective continuous functions map $\mathcal{E}$ into $\mathcal{E}$.*

[4] In the literature, the effective continuous functions on $\mathcal{P}(\omega)$ are called **enumeration operators** (see Odifreddi [1999], Chapter XIV). They constitute an analogue of the partial recursive functions on ω, which are defined by the similar condition that their graph is an r.e. set.

Proof. Let f be an effective continuous function, and A be an r.e. set. Recall that $f = I(J(f))$, because $I \circ J$ is the identity by 13.3.2. Then

$$x \in f(A) \iff (\exists n)(\langle x, n\rangle \in J(f) \ \wedge \ u_n \subseteq A)$$

by definition of I. Since $J(f)$ and A are r.e. by hypothesis, and the r.e. sets are closed under existential quantification, the right-hand-side is r.e. □

The previous proposition allows us to restrict attention from arbitrary sets of natural numbers and continuous functions, to r.e. sets and effective continuous functions. The next result is the crucial verification needed to prove that, by so doing, we still obtain a model of the Untyped Lambda Calculus.

Proposition 13.3.6 (Plotkin [1972], Scott [1975]) *For any environment ρ on $\mathcal{E}$ and any term t, $[\![t]\!]_\rho \in \mathcal{E}$.*

Proof. Recall that

$$[\![t]\!]_\rho = \begin{cases} \rho(x) & \text{if } t = x \\ I([\![u]\!]_\rho)([\![v]\!]_\rho) & \text{if } t = uv \\ J(\Lambda X.\, [\![u]\!]_{\rho[x:=X]}) & \text{if } t = \lambda x.\, u, \end{cases}$$

where $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ denotes the function

$$A \in \mathcal{E} \longmapsto [\![u]\!]_{\rho[x:=A]} \in \mathcal{E}.$$

We proceed by induction on the definition of $[\![t]\!]_\rho$:

1. $[\![x]\!]_\rho = \rho(x)$ is an r.e. set by the hypothesis on ρ.

2. Since an effective continuous function maps r.e. sets to r.e. sets, to show inductively that $[\![uv]\!]_\rho = I([\![u]\!]_\rho)([\![v]\!]_\rho)$ is an r.e. set it is enough to prove that $I(X)(Y)$ is an effective continuous function of X and Y. This is immediate by definition, since

   $$x \in I(X)(Y) \ \Leftrightarrow \ (\exists n)(\langle x, n\rangle \in X \wedge u_n \subseteq Y),$$

   and the right-hand-side is r.e.

3. Since the graph of an effective continuous function is an r.e. set, to show that $[\![\lambda x.\, u]\!]_\rho = J(\Lambda X.\, [\![u]\!]_{\rho[x:=X]})$ is an r.e. set it is enough to prove that $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ is an effective continuous function. This can be done inductively on u:

   - If $u = x$, then $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ is the identity function $\Lambda X.\, X$, which is an effective continuous function.

- If $u = y \neq x$, then $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ is the constant function $\Lambda X.\, \rho(y)$, which is an effective continuous function.
- If $u = u_1 u_2$, then $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ is $\Lambda X.\, I([\![u_1]\!]_{\rho[x:=X]})([\![u_2]\!]_{\rho[x:=X]})$, which is an effective continuous function by the induction hypothesis, closure under composition and the fact, proved in part 2, that $I(X)(Y)$ is an effective continuous function of X and Y.
- If $u = \lambda y.\, u_1$, then $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ is $\Lambda X.\, J(\Lambda Y.\, [\![u_1]\!]_{\rho[x:=X;y:=Y]})$, which is an effective continuous function by the induction hypothesis, closure under composition and the fact that if $f(X,Y)$ is an effective continuous function of X and Y, then $J(\Lambda Y.\, f(X,Y))$ is an effective continuous function of X, since

$$\langle z, n\rangle \in J(\Lambda Y.\, f(X,Y)) \ \Leftrightarrow\ z \in f(X, u_n),$$

and the right-hand-side is r.e. □

Actually, for any environment ρ on $\mathcal{P}(\omega)$ and any closed term t we have $[\![t]\!]_\rho \in \mathcal{E}$.

Corollary 13.3.7 The Effective Graph Model (Plotkin [1972], Scott [1975]) *$\mathcal{E}$ is a (nonextensional) countable model of the Untyped Lambda Calculus.*

Proof. By the proof of 13.3.2 and the previous proposition. □

13.4 Inverse Limits

We now look for extensional models of Lambda Calculus. By Condition 2 of 13.2.2, it would be enough to go one step beyond the work done so far, and find a c.c.p.o. D and continuous functions

$$I : D \to [D \to D] \qquad \text{and} \qquad J : [D \to D] \to D$$

such that not only $I \circ J$ is the identity on $[D \to D]$, but also $J \circ I$ is the identity on D. In other words, we would like to find a c.c.p.o. D isomorphic to its own function space $[D \to D]$.

The idea to obtain such a D is quite simple. If we identify two c.c.p.o.'s when there are continuous functions between them whose compositions are the identities, and let

$$F(X) = [X \to X],$$

then such a D is a fixed point of F, i.e. $D = F(D)$. We now first discover a condition ensuring the existence of fixed points, and then show that F satisfies the condition.

Existence of fixed points

The following simple but powerful result exhibits a sufficient condition for a function f to have a fixed point.

Proposition 13.4.1 (Knaster [1928], Tarski [1955], Abian and Brown [1961]) *If $(D, \sqsubseteq)$ is a c.c.p.o. and $f : D \to D$ is a continuous function, then f has a least fixed point.*

Moreover, if f is expansionary, i.e. $x \sqsubseteq f(x)$ for every x, then f has a fixed point above any given $x \in D$.

Proof. We iterate f starting from $\perp$, as follows:

$$\begin{array}{lcl} x_0 & = & \perp \\ x_{n+1} & = & f(x_n) \\ x_\infty & = & \bigsqcup_{n\in\omega} x_n. \end{array}$$

First, we show that x_∞ *exists*. For this it is enough to show that the set $\{x_n\}_{n\in\omega}$ is actually a chain, and thus it has a l.u.b. This is proved by induction:

- $x_0 \sqsubseteq x_1$ because $x_0 = \perp$
- if $x_n \sqsubseteq x_{n+1}$, then
$$x_{n+1} = f(x_n) \sqsubseteq f(x_{n+1}) = x_{n+2}$$
by monotonicity of f, which is a continuous function.

Second, we show that x_∞ is a *fixed point* of f. Indeed,

$$f(x_\infty) = f(\bigsqcup_{n\in\omega} x_n) = \bigsqcup_{n\in\omega} f(x_n) = \bigsqcup_{n\in\omega} x_{n+1} = x_\infty$$

by definition of x_∞, continuity of f, definition of x_{n+1}, and the fact that the chain $\{x_{n+1}\}_{n\in\omega}$ has l.u.b. x_∞ (because it is obtained from the chain $\{x_n\}_{n\in\omega}$ by dropping the first element, and this has no influence on the l.u.b.).

Finally, we show that x_∞ is the *least* fixed point of f. Let a be any fixed point, i.e. $f(a) = a$. By the properties of l.u.b.'s, to show that $x_\infty \sqsubseteq a$ it is enough to show that, for every n, $x_n \sqsubseteq a$. This is immediate by induction:

- $x_0 \sqsubseteq a$ because $x_0 = \perp$
- if $x_n \sqsubseteq a$, then
$$x_{n+1} = f(x_n) \sqsubseteq f(a) = a$$
by definition of x_{n+1}, monotonicity of f, and the fact that a is a fixed point.

If f is *expansionary* then we can proceed as above, by starting from $x_0 = x$ for any given x. Since f is expansionary, $x_0 \sqsubseteq f(x_0) = x_1$. Then, as above, x_∞ exists and it is a fixed point of f. And, by definition, $x = x_0 \sqsubseteq x_\infty$. □

Exercises 13.4.2 More on fixed points. (Knaster [1928], Tarski [1955], Abian and Brown [1961])

a) *A monotone function f on a complete lattice L has a least fixed point and a greatest fixed point.* (Hint: for the least fixed point, consider the set $\{x : f(x) \sqsubseteq x\}$, which is not empty because 1 belongs to it. This set contains every fixed point of f, and its g.l.b. is the least fixed point of f. Similarly, the l.u.b. of the set $\{x : f(x) \sqsupseteq x\}$ is the greatest fixed point of f.)

b) *The set of fixed points of a monotone function f on a complete lattice L is itself a complete lattice.* (Hint: since the l.u.b. in L of a set of fixed points of f is not necessarily a fixed point, consider the least fixed point greater than it. Similarly for the g.l.b.)

c) *A monotone function f on a c.c.p.o. D has a least fixed point.* (Hint: the proof of 13.4.1 can be extended through the ordinals, by letting $x_{\alpha+1} = f(x_\alpha)$, and $x_\alpha = \bigsqcup_{\beta<\alpha} x_\beta$ if α is a limit ordinal. Since this is a chain of elements of D, its length cannot exceed the maximal length of such chains, and hence it reaches a fixed point.)

d) *An expansionary function f on a c.c.p.o. has a fixed point.* (Hint: the construction of part c) still produces a chain of elements, and hence a fixed point, although not necessarily a least one.)

Methodologically, part a) uses a stronger hypothesis on the domain and it constructs the least fixed point from above, by a purely algebraic proof. Part c) uses instead a weaker hypothesis on the domain and it constructs the least fixed point from below, by a set-theoretical proof using ordinals.

Moreover, part d) implies Zorn's Lemma in the form: *any c.c.p.o. D has a maximal element*. Otherwise $(\forall x \in D)(\exists y \in D)(x \sqsubseteq y \wedge x \neq y)$, and there is a choice function f on D such that $(\forall x \in D)(x \sqsubseteq f(x) \wedge x \neq f(x))$. Then f is expansionary and has no fixed point.

Exercises 13.4.3 Uniform fixed point operators. A **fixed point operator** is a family $\mathbf{F} = \{\mathbf{F}_D\}_D$ such that, for any c.c.p.o. D and any continuous function f on D, $\mathbf{F}_D(f)$ is a fixed point of f.

A fixed point operator is **uniform** if, for any c.c.p.o. D_1 and D_2 and any function $f_1 \in [D_1 \to D_1]$, $f_2 \in [D_2 \to D_2]$ and $g : D_1 \to D_2$ such that $g(\perp_{D_1}) = \perp_{D_2}$ and $g \circ f_1 = f_2 \circ g$, then $g(\mathbf{F}_{D_1}(f_1)) = \mathbf{F}_{D_2}(f_2)$.

a) *The least fixed point operator* $\mathbf{Fix} = \{\mathbf{Fix}_D\}_D$ *is a uniform fixed point operator.* (Hint: notice that

$$g(f_1^{(n)}(\perp_{D_1})) = f_2^{(n)}(g(\perp_{D_1})) = f_2^{(n)}(\perp_{D_2})$$

by the hypotheses on f_1, f_2 and g. Then

$$g(\mathbf{Fix}_{D_1}(f_1)) = g(\bigsqcup f_1^{(n)}(\perp_{d_1})) = \bigsqcup g(f_1^{(n)}(\perp_{D_1})) = \bigsqcup f_2^{(n)}(\perp_{D_2}) = \mathbf{Fix}_{D_2}(f_2)$$

by the construction of the least fixed point in 13.4.1, and continuity of g.)

b) *The least fixed point operator* $\mathbf{Fix} = \{\mathbf{Fix}_D\}_D$ *is the only uniform fixed point operator.* (Hint: suppose $\mathbf{F}$ is a uniform fixed point operator. Given D and $f \in [D \to D]$, consider the restriction D_1 of D below the least fixed point of f, and the restriction f_1 of f to D_1. Then $f_1 \in [D_1 \to D_1]$, because

$$x \in D_1 \;\Rightarrow\; x \sqsubseteq \mathbf{Fix}_D(f) \;\Rightarrow\; f(x) \sqsubseteq f(\mathbf{Fix}_D(f)) = \mathbf{Fix}_D(f) \;\Rightarrow\; f(x) \in D_1.$$

Moreover, the least fixed point of f is the only fixed point of f_1.

If g is the identity embedding of D_1 into D, then

$$g(\mathbf{F}_{D_1}(f_1)) = \mathbf{F}_D(f)$$

because $\mathbf{F}$ is uniform. And

$$g(\mathbf{F}_{D_1}(f_1)) = \mathbf{F}_{D_1}(f_1) = \mathbf{Fix}_D(f)$$

because f_1 has a unique fixed point on D_1. Then $\mathbf{F}_D = \mathbf{Fix}_D$ and $\mathbf{F} = \mathbf{Fix}$.)

Plan of the proof

We now have a general condition for the existence of fixed points. The strategy to show that $F(X) = [X \to X]$ has a fixed point is then to show that it is continuous. But F is a function whose arguments and values are c.c.p.o.'s, and we do not know what continuity means for F. The general plan toward our goal is thus the following:

1. consider the collection $\mathcal{D}$ of all c.c.p.o.'s

2. define a 'partial ordering' $\trianglelefteq$ on $\mathcal{D}$, among c.c.p.o.'s

3. show that $(\mathcal{D}, \trianglelefteq)$ is a 'c.c.p.o.', i.e. every chain of elements of $\mathcal{D}$ has a l.u.b.

4. to have fixed points, show that F is 'continuous' as a function on $\mathcal{D}$, i.e. it preserves l.u.b.'s of chains

5. to have fixed points above any given element, show that F is expansionary as a function on $\mathcal{D}$.

We say that $\trianglelefteq$ is a 'partial ordering' (in quotes) because it is such not on the elements of $\mathcal{D}$ directly, but only on equivalence classes of them. We say that $(\mathcal{D}, \trianglelefteq)$ is a 'c.c.p.o.' because $\mathcal{D}$ is too big to be a real c.c.p.o. (in technical terms: it is not a set, but a class). We say that F is 'continuous' because we will only prove that *special* chains have l.u.b.'s. Then we will not be able to appeal directly to 13.4.1 above, but a reproduction of its proof will produce the desired result.

Partial ordering among c.c.p.o.'s

We now face our first task, of defining a partial ordering $\trianglelefteq$ on the collection $\mathcal{D}$ of all c.c.p.o.'s. The idea is that $D \trianglelefteq D'$ means that D' has more information than D. Thus we should be able to identify all elements of D with corresponding elements of D', although not necessarily the converse.

Definition 13.4.4 *Given two c.c.p.o.'s $(D, \sqsubseteq_D)$ and $(D', \sqsubseteq_{D'})$, we say that $\boldsymbol{D \trianglelefteq D'}$ if there are two continuous functions*

$$i : D \longrightarrow D' \qquad \textit{and} \qquad j : D' \longrightarrow D$$

such that, for all $x \in D$ and $y \in D'$,

$$j(i(x)) = x \qquad \textit{and} \qquad i(j(y)) \sqsubseteq_{D'} y.$$

The intuition is that $i(x)$ is a copy of x in D', while $j(y)$ is only the best approximation of y in D. Then the first condition tells that x is the best approximation in D of its own copy $i(x)$ in D', and the second condition tells that when approximating y in D we may lose some information.

Notice that the first condition implies that *i is one-one and j is onto*. Oneoneness follows from the fact that

$$i(x_1) = i(x_2) \;\Rightarrow\; j(i(x_1)) = j(i(x_2)) \;\Rightarrow\; x_1 = x_2.$$

Ontoness follows from the fact that $i(x)$ a counterimage of $x \in D$ under j, since $j(i(x)) = x$.

Moreover, *i and j are adjoint* in the sense of 5.1.3, i.e.

$$i(x) \sqsubseteq_{D'} y \iff x \sqsubseteq_D j(y).$$

If $i(x) \sqsubseteq y$, then $x = j(i(x)) \sqsubseteq j(y)$ by the hypothesis on $j \circ i$ and monotonicity of j. And if $x \sqsubseteq j(y)$, then $i(x) \sqsubseteq i(j(y)) \sqsubseteq y$ by monotonicity of i and the hypothesis on $i \circ j$.

Notice also that *$\trianglelefteq$ is a reflexive and transitive relation*:

1. $D \trianglelefteq D$ via the identity function (in both directions)
2. if $D_1 \trianglelefteq D_2$ via i_1 and j_1, and $D_2 \trianglelefteq D_3$ via i_2 and j_2, then $D_1 \trianglelefteq D_3$ via $i_2 \circ i_1$ and $j_1 \circ j_2$.

This shows that $\trianglelefteq$ is 'almost' a partial ordering. The only missing property is antisymmetry, which does not hold because two c.c.p.o.'s may be isomorphic without being equal. But this is easily taken care of by modifying the notion of equality, and adapting it to our needs.

Definition 13.4.5 *Given two c.c.p.o.'s D_1 and D_2, we let* $\boldsymbol{D_1 = D_2}$ *if and only if $D_1 \trianglelefteq D_2$ and $D_2 \trianglelefteq D_1$.*

Then $=$ is an equivalence relation, and $\trianglelefteq$ induces a partial ordering on the equivalence classes. The least c.c.p.o. w.r.t. to the partial ordering induced by $\trianglelefteq$ is the trivial c.c.p.o. $\{\perp\}$, consisting of only one element.

Least upper bounds of chains of c.c.p.o.'s

We now turn to our second task, of showing that chains of c.c.p.o.'s have l.u.b. Given c.c.p.o.'s $(D_n, \sqsubseteq_n)$ such that

$$D_0 \trianglelefteq D_1 \trianglelefteq \cdots \trianglelefteq D_n \trianglelefteq D_{n+1} \trianglelefteq \cdots,$$

we want a c.c.p.o. $(D_\infty, \sqsubseteq_\infty)$ such that:

- D_∞ is an *upper bound* of $\{D_n\}_{n\in\omega}$, i.e. $D_n \trianglelefteq D_\infty$
- D_∞ is the *least* upper bound, i.e. given any D such that $D_n \trianglelefteq D$ for every n, then $D_\infty \trianglelefteq D$.

Definition 13.4.6 *Given a chain*

$$D_0 \trianglelefteq D_1 \trianglelefteq \cdots \trianglelefteq D_n \trianglelefteq D_{n+1} \trianglelefteq \cdots$$

of c.c.p.o.'s $(D_n, \sqsubseteq_n)$, with continuous functions

$$i_n : D_n \longrightarrow D_{n+1} \qquad \textit{and} \qquad j_n : D_{n+1} \longrightarrow D_n$$

such that, for all $x_n \in D_n$ and $x_{n+1} \in D_{n+1}$,

$$j_n(i_n(x_n)) = x_n \qquad \textit{and} \qquad i_n(j_n(x_{n+1})) \sqsubseteq_{n+1} x_{n+1},$$

we define a c.c.p.o. $\boldsymbol{(D_\infty, \sqsubseteq_\infty)}$ *as follows:*

1. *D_∞ is the set of all sequences $x_\infty = \langle x_n\rangle_{n\in\omega}$ such that*

$$(\forall n)(x_n \in D_n) \qquad \textit{and} \qquad (\forall n)(j_n(x_{n+1}) = x_n)$$

2. *$\sqsubseteq_\infty$ is the order componentwise, i.e.*

$$\langle x_n\rangle_{n\in\omega} \sqsubseteq_\infty \langle y_n\rangle_{n\in\omega} \;\Leftrightarrow\; (\forall n)(x_n \sqsubseteq_n y_n).$$

The idea underlying the definition of D_∞ is quite simple, and it mimicks one possible construction of the real numbers. Think of D_n as the real numbers with decimal expansion truncated at the n-th digit, and of D_∞ as the reals. There are natural choices for i_n and j_n. Namely, i_n is the function that adds a 0 at the end of any element of D_n, thus transforming it into an equivalent element of D_{n+1}. And j_n is the function that forgets the last digit of an element of D_{n+1}, thus losing some information. Then an element of D_∞ can be seen as an infinite sequence of truncated approximations, each forgetting the last digit of the following one. This is the reason for the condition using the j_n's, which also gives the name of **inverse limit** to the construction. (Each approximation also adds a new digit to the previous one, but not necessarily a 0, which is why we do not use the i_n's instead).

It is easy to see that $(D_\infty, \sqsubseteq_\infty)$ *is a c.c.p.o.* Precisely, if

$$x_\infty^0 \sqsubseteq_\infty x_\infty^1 \sqsubseteq_\infty \cdots$$

where $x_\infty^m = \langle x_n^m \rangle_{n\in\omega}$, then

$$\bigsqcup_{m\in\omega} x_\infty^m = \bigsqcup_{m\in\omega} \langle x_n^m \rangle_{n\in\omega} = \langle \bigsqcup_{m\in\omega} x_n^m \rangle_{n\in\omega}.$$

The first equality holds by definition of x_n^∞. The second is immediate, once we recall that the ordering on D_∞ is defined componentwise and that each D_n is a c.c.p.o.

Notice that $\perp_{D_\infty} = \langle \perp_{D_n} \rangle_{n\in\omega}$.

We now have to prove that $(D_\infty, \sqsubseteq_\infty)$ really is the least upper bound of $\{D_n\}_{n\in\omega}$.

Proposition 13.4.7 *D_∞ is an upper bound to the chain $\{D_n\}_{n\in\omega}$, i.e.*

$$D_n \trianglelefteq D_\infty$$

for every n.

Proof. We need continuous functions

$$i_{n,\infty} : D_n \longrightarrow D_\infty \qquad \text{and} \qquad j_{n,\infty} : D_\infty \longrightarrow D_n$$

such that, for all $x_n \in D_n$ and $x_\infty \in D_\infty$,

$$j_{n,\infty}(i_{n,\infty}(x_n)) = x_n \qquad \text{and} \qquad i_{n,\infty}(j_{n,\infty}(x_\infty)) \sqsubseteq_\infty x_\infty.$$

The idea behind the definitions of $i_{n,\infty}$ and $j_{n,\infty}$ is obvious. For $i_{n,\infty}$, we embed x_n in D_∞ in the natural way, by considering a sequence of approximations whose first n elements are the approximations of x_n at the previous levels, and the remaining elements are successive copies of x_n at the following levels. For $j_{n,\infty}$, we simply project an element of D_∞ at the appropriate level. Formally, we let

$$i_{n,\infty}(x_n) = \langle \ldots, j_{n-2}(j_{n-1}(x_n)), j_{n-1}(x_n), x_n, i_n(x_n), i_{n+1}(i_n(x_n)), \ldots \rangle$$

and

$$j_{n,\infty}(\langle x_n \rangle_{n \in \omega}) = x_n.$$

Clearly $i_{n,\infty}(x_n) \in D_\infty$, because each component is at the appropriate level by definition, and is obtained from the following one by the j functions. This is so by definition for the first n levels, and by the fact that $j_m(i_m(x_m)) = x_m$ for the remaining ones, e.g.

$$j_{n+1}(i_{n+1}(i_n(x_n))) = i_n(x_n).$$

Similarly, $j_{n,\infty}(x_\infty) \in D_n$ by definition of D_∞.

Since x_n *is* the n-th component of $i_{n,\infty}(x_n)$, we have

$$j_{n,\infty}(i_{n,\infty}(x_n)) = x_n. \tag{13.1}$$

To show

$$i_{n,\infty}(j_{n,\infty}(x_\infty)) \sqsubseteq_\infty x_\infty \tag{13.2}$$

we proceed componentwise, by the definition of $\sqsubseteq_\infty$. Notice that

$$\begin{array}{rcl} i_{n,\infty}(j_{n,\infty}(x_\infty)) & = & \langle \ldots, j_{n-1}(x_n), x_n, i_n(x_n), \ldots \rangle \\ x_\infty & = & \langle \ldots\ldots, x_{n-1}, x_n, x_{n+1}, \ldots\ \rangle \end{array}$$

where, for every m, $x_m = j_m(x_{m+1})$. The first n components are thus equal. We then proceed by induction:

- Since $x_n = j_n(x_{n+1})$,

$$i_n(x_n) = i_n(j_n(x_{n+1})) \sqsubseteq_{n+1} x_{n+1}.$$

- From $i_n(x_n) \sqsubseteq_{n+1} x_{n+1} = j_{n+1}(x_{n+2})$, we obtain

$$i_{n+1}(i_n(x_n)) \sqsubseteq_{n+2} i_{n+1}(j_{n+1}(x_{n+2})) \sqsubseteq_{n+2} x_{n+2}$$

by monotonicity of i_{n+1}, and so on.

We leave to the reader the verification that $i_{n,\infty}$ and $j_{n,\infty}$ are continuous. □

In terms of the previous analogy with real numbers, we can think of $i_{n,\infty}$ as the function that adds infinitely many 0's to any element of D_n, and of $j_{n,\infty}$ as the function that truncates the infinite decimal expansion of a real number at the

Given c.c.p.o.'s $(D_n, \sqsubseteq_n)$ such that

$$D_0 \trianglelefteq \cdots \trianglelefteq D_n \trianglelefteq \cdots,$$

with continuous functions

$$i_n : D_n \to D_{n+1} \qquad \text{and} \qquad j_n : D_{n+1} \to D_n$$

such that, for every n and every $x_n \in D_n$,

$$j_n(i_n(x_n)) = x_n \qquad \text{and} \qquad i_n(j_n(x_{n+1})) \sqsubseteq_{n+1} x_{n+1},$$

we define:

$$\begin{array}{rcl}
\langle x_n \rangle_{n\in\omega} \in D_\infty & \Leftrightarrow & (\forall n)(x_n \in D_n \ \wedge \ x_n = j_n(x_{n+1})) \\
\langle x_n \rangle_{n\in\omega} \sqsubseteq_\infty \langle y_n \rangle_{n\in\omega} & \Leftrightarrow & (\forall n)(x_n \sqsubseteq_n y_n) \\
i_{n,\infty}(x_n) & = & \langle \ldots, j_{n-1}(x_n), x_n, i_n(x_n), \ldots \rangle \\
j_{n,\infty}(\langle x_n \rangle_{n\in\omega}) & = & x_n,
\end{array}$$

where

$$i_{n,\infty} : D_n \to D_\infty \qquad \text{and} \qquad j_{n,\infty} : D_\infty \to D_n.$$

Then, for every $x_\infty = \langle x_n \rangle_{n\in\omega} \in D_\infty$,

$$\langle x_n \rangle_{n\in\omega} = \bigsqcup_{n\in\omega} i_{n,\infty}(x_n) \qquad \text{and} \qquad x_\infty = \langle j_{n,\infty}(x_\infty) \rangle_{n\in\omega}.$$

Moreover,

1. $j_{n,\infty}(i_{n,\infty}(x_n)) = x_n$
2. $i_{n,\infty}(j_{n,\infty}(x_\infty) \sqsubseteq_\infty x_\infty$
3. $i_{n,\infty}(x_n) = i_{n+1,\infty}(i_n(x_n))$
4. $j_{n,\infty}(x_\infty) = j_n(j_{n+1,\infty}(x_\infty))$
5. $\bigsqcup_{n\in\infty} i_{n,\infty}(j_{n,\infty}(x_\infty)) = x_\infty$.

Figure 13.1: Definition and properties of D_∞

n-th digit. Thus they really are infinitary versions of the finitary i_n and j_n, that work one level at a time.

Notice the following interesting amalgamation properties, relating the finitary (i_n and j_n) and infinitary ($i_{n,\infty}$ and $j_{n,\infty}$) embedding and projection functions defined above:

$$i_{n,\infty}(x_n) \quad = \quad i_{n+1,\infty}(i_n(x_n)) \tag{13.3}$$

and

$$j_{n,\infty}(\langle x_m\rangle_{m\in\omega}) \quad = \quad j_n(j_{n+1,\infty}(\langle x_m\rangle_{m\in\omega})). \tag{13.4}$$

Intuitively, they say that we can embed an element x_n of level n into D_∞ either directly, or by first embedding it one level up, and then embedding the result into D_∞. Similarly, we can project an element x_∞ of D_∞ at level n either directly, or by first projecting it at level $n+1$, and then projecting the result one level down.

To prove 13.3, we notice that the left-hand-side is a sequence with x_n in the n-th position, that proceeds on the right by using i's and on the left by using j's. The right-hand-side is a sequence with $i_n(x_n)$ in the $n+1$-th position, that proceeds as the previous one. The two sequences are equal, because $x_n = j_n(i_n(x_n))$.

To prove 13.4, we notice that the left-hand-side is x_n, the right-hand-side is $j_n(x_{n+1})$, and they are equal as above.

We also notice that, if $x_\infty = \langle x_n\rangle_{n\in\omega} \in D_\infty$,

$$\langle x_n\rangle_{n\in\omega} = \bigsqcup_{n\in\omega} i_{n,\infty}(x_n) \quad \text{and} \quad x_\infty = \langle j_{n,\infty}(x_\infty)\rangle_{n\in\omega}. \tag{13.5}$$

Indeed, $i_{n,\infty}(x_n)$ is a sequence that coincides with x_∞ in the first n places, by definition. Then the n-th place of these sequences is eventually constant and equal to x_n, and the l.u.b. of all these sequences is $\langle x_n\rangle_{n\in\omega}$ (recall that the order $\sqsubseteq_\infty$ is defined componentwise).

Similarly, $j_{n,\infty}(x_\infty)$ is x_n, and thus the sequence $\langle j_{n,\infty}(x_\infty)\rangle_{n\in\omega}$ is actually x_∞ itself.

13.5 vindicates the intuition that led us to the definition of D_∞, when we thought of the elements x_n as approximations of $\langle x_n\rangle_{n\in\omega}$. Now this can be made precise, in two complementary senses. First, given a sequence of approximations, we can identify it with the l.u.b. of the set of elements obtained by embedding each approximation into D_∞. Second, given an element of D_∞, we can identify it with the sequence of its own approximations. This allows us to use, in any situation, the most convenient of the two ways of seeing an element of D_∞: either as a completed process of approximations, hence as a single limit *element*, or as the process of approximation itself, hence as an infinite *sequence*. The two ways are related to two different conceptions of infinity, namely actual and potential.

Figure 13.1 collects all facts about D_∞ proved so far.

Proposition 13.4.8 *D_∞ is an 'almost' least upper bound to the chain $\{D_n\}_{n\in\omega}$. Precisely, it is the least among the upper bounds D such that, for every n, there are continuous functions*

$$f_n : D_n \longrightarrow D \qquad \textit{and} \qquad g_n : D \longrightarrow D_n$$

satisfying, for all $x_n \in D_n$ and $x \in D$, not only

1. $g_n(f_n(x_n)) = x_n$
2. $f_n(g_n(x)) \sqsubseteq_D x$,

as for a usual upper bound, but also

3. $f_n(x_n) = f_{n+1}(i_n(x_n))$
4. $g_n(x) = j_n(g_{n+1}(x))$.

Under these conditions,

$$D_\infty \trianglelefteq D.$$

Proof. First, notice that D_∞ is such an upper bound, since we can let $f_n = i_{n,\infty}$ and $g_n = j_{n,\infty}$. Then Conditions 1 and 2 are satisfied by the previous result, as they are for any upper bound, and Conditions 3 and 4 were noticed in 13.3 and 13.4 just above.

Given D as stated, we now prove that $D_\infty \trianglelefteq D$. We need continuous functions

$$f_\infty : D_\infty \to D \qquad \text{and} \qquad g_\infty : D \to D_\infty$$

such that, for all $x_\infty \in D_\infty$ and $x \in D$,

$$g_\infty(f_\infty(x_\infty)) = x_\infty \qquad \text{and} \qquad f_\infty(g_\infty(x)) \sqsubseteq_D x.$$

Given an element $x_\infty = \langle x_n \rangle_{n\in\omega}$ of D_∞, f_∞ must associate to it an element of D. Since x_n is an approximation to x_∞, and $f_n(x_n)$ is an element of D, it is natural to consider the latter as an approximation of the needed element, and to define

$$f_\infty(\langle x_n \rangle_{n\in\omega}) = \bigsqcup_{n\in\omega} f_n(x_n).$$

Such an element exists if $\{f_n(x_n)\}_{n\in\omega}$ is a chain, because D is a c.c.p.o. Notice that

$$i_n(x_n) = i_n(j_n(x_{n+1})) \sqsubseteq_{n+1} x_{n+1},$$

because $\langle x_n \rangle_{n\in\omega} \in D_\infty$, i.e. $x_n = j_n(x_{n+1})$, and by properties of i_n and j_n. Then, by monotonicity of f_{n+1},

$$f_{n+1}(i_n(x_n)) \sqsubseteq_D f_{n+1}(x_{n+1}).$$

We can then claim that $f_n(x_n) \sqsubseteq_D f_{n+1}(x_{n+1})$ if

$$f_n(x_n) = f_{n+1}(i_n(x_n)),$$

and this is precisely the additional Condition 3 postulated above.[5]

Given an element $x \in D$, g_∞ must associate to it an element of D_∞, hence a sequence of elements in the various D_n. The g_n provide such elements, and it is thus natural to let

$$g_\infty(x) = \langle g_n(x) \rangle_{n \in \omega}.$$

Such a sequence is in D_∞, if its elements are related as follows:

$$g_n(x) = j_n(g_{n+1}(x)),$$

and this is precisely the additional Condition 4 postulated above.

This shows that if we let

$$f_\infty(\langle x_n \rangle_{n\in\omega}) = \bigsqcup_D f_n(x_n) \qquad \text{and} \qquad g_\infty(x) = \langle g_n(x) \rangle_{n\in\omega},$$

we obtain functions from D_∞ to D and conversely. We now show that they have the needed properties.

1. *for any $x \in D$, $f_\infty(g_\infty(x)) \sqsubseteq_D x$*
   This is easily checked:
   $$\begin{array}{rcl} f_\infty(g_\infty(x)) & = & f_\infty(\langle g_n(x) \rangle_{n\in\omega}) \\ & = & \bigsqcup_{n\in\omega} f_n(g_n(x)) \\ & \sqsubseteq_D & x \end{array}$$
   by the definitions of g_∞ and f_∞, and the fact that $f_n(g_n(x)) \sqsubseteq_D x$ for every n.

2. *for any $x_\infty = \langle x_n \rangle_{n\in\omega} \in D_\infty$, $g_\infty(f_\infty(x_\infty)) = x_\infty$*
   This is the only nontrivial verification:
   $$\begin{array}{rcl} g_\infty(f_\infty(x_\infty)) & = & g_\infty(\bigsqcup\limits_{m\in\omega} f_m(x_m)) \\ & = & \langle g_n(\bigsqcup\limits_{m\in\omega} f_m(x_m)) \rangle_{n\in\omega} \\ & = & \langle \bigsqcup\limits_{m\in\omega} g_n(f_m(x_m)) \rangle_{n\in\omega} \\ & = & \langle x_n \rangle_{n\in\omega} \\ & = & x_\infty \end{array}$$

[5] Notice that the weaker condition $f_n(x_n) \sqsubseteq_D f_{n+1}(i_n(x_n))$ would actually suffice, but we will have no use of this added generality in the cases we are interested in.

by the definitions of f_∞ and g_∞, continuity of g_n, property 13.6 proved below, and the definition of x_∞.

To complete the proof of part 2 above it remains to check that, for every n,

$$\bigsqcup_{m\in\omega} g_n(f_m(x_m)) = x_n. \tag{13.6}$$

There are two cases:

- *if $m \leq n$, then $g_n(f_m(x_m)) \sqsubseteq_n x_n$*
  If $m \leq n$, then $f_m(x_m) \sqsubseteq_D f_n(x_n)$ by what proved above. Then

  $$g_n(f_m(x_m)) \sqsubseteq_n g_n(f_n(x_n)) = x_n$$

  by monotonicity of g_n and Condition 1.

- *if $m \geq n$, then $g_n(f_m(x_m)) = x_n$*
  This is proved by starting with $g_m(f_m(x_m)) = x_m$, and successively using projections on both sides, together with the properties

  $$j_m(x_{m+1}) = x_m \qquad \text{and} \qquad j_m(g_{m+1}(x)) = g_m(x).$$

  For example,

  $$\begin{array}{rrcl} & g_{n+1}(f_{n+1}(x_{n+1})) & = & x_{n+1} \\ \Rightarrow & j_n(g_{n+1}(f_{n+1}(x_{n+1}))) & = & j_n(x_{n+1}) \\ \Rightarrow & g_n(f_{n+1}(x_{n+1})) & = & x_n. \end{array}$$

We leave to the reader the verification that f_∞ and g_∞ are continuous, using the fact that f_n and g_n are continuous, for all n. □

We have proved that *for any given chain $\{D_n\}_{n\in\omega}$, D_∞ is uniquely determined up to equality of c.c.p.o.'s*. Notice that *if the following additional condition is satisfied:*

5. $\bigsqcup_{n\in\omega} f_n(g_n(x)) = x$,

then D is actually equal to D_∞, because in part 1 of the above proof we can conclude

$$f_\infty(g_\infty(x)) = x.$$

Since the symmetrical condition

$$g_\infty(f_\infty(x_\infty)) = x_\infty$$

always holds, it follows that f_∞ and g_∞ are inverse one of the other. In particular, they are both one-one and onto, and preserve structure by continuity. We have thus obtained a useful *sufficient condition to show that D is isomorphic to D_∞*.

In the special case of $D = D_\infty$, where $f_n = i_{n,\infty}$ and $g_n = j_{n,\infty}$, Condition 5 becomes

$$\bigsqcup_{n\in\omega} i_{n,\infty}(j_{n,\infty}(x_\infty)) = x_\infty$$

and is obviously satisfied (see 13.5), while the functions f_∞ and g_∞ defined above reduce to

$$i_\infty(\langle x_n\rangle_{n\in\omega}) = \bigsqcup_{n\in\omega} i_{n,\infty}(x_n) \qquad \text{and} \qquad j_\infty(x_\infty) = \langle j_{n,\infty}(x_\infty)\rangle_{n\in\omega},$$

which are actually both the identity function on D_∞ (see 13.5).

Continuity of the function space operator

$\mathcal{D}$ fails being a c.c.p.o. in two different senses. First, because its domain is not a set. Second, because 13.4.8 provides only special upper bounds for chains, not necessarily *least* ones. Then we cannot directly apply 13.4.1 on the existence of fixed points, but we can certainly redo its proof, and check that it works.

The reason why we stressed the existence of fixed point for f *above* any given element, when f is expansionary, is that the least fixed point of $F(X) = [X \to X]$ is trivial, namely the one-element c.c.p.o. $D = \{\perp\}$. Indeed, there is only one function from D to D, and such a function is continuous. Thus $[D \to D]$ is a one-element c.c.p.o. isomorphic to D. But we are looking for *nontrivial* solutions, and they can be obtained if F has a fixed point above any given c.c.p.o. D.

For the rest of this section, F is the function on $\mathcal{D}$ defined by

$$F(X) = [X \to X].$$

Proposition 13.4.9 *F is expansionary. Precisely, for every D,*

$$D \trianglelefteq [D \to D].$$

Proof. We prove that the functions

$$i : D \to [D \to D] \qquad \text{and} \qquad j : [D \to D] \to D,$$

naturally defined by

$$\begin{aligned} i(x) &= \text{the constant function on } D \text{ with value } x \\ j(f) &= f(\perp_D), \end{aligned}$$

are continuous, and

$$j(i(x)) = x \qquad \text{and} \qquad i(j(f)) \sqsubseteq_{[D\to D]} f$$

for all $x \in D$ and $f \in [D \to D]$.

The first fact is trivially true, since $j(i(x))$ is a particular value of the constant function $i(x)$ with value x.

By definition of $\sqsubseteq_{[D \to D]}$, to prove $i(j(f)) \sqsubseteq_{[D \to D]} f$ we need to prove

$$i(j(f))(x) \sqsubseteq_D f(x)$$

for every $x \in D$. But

$$i(j(f))(x) = i(f(\perp_D))(x) = f(\perp_D) \sqsubseteq_D f(x)$$

by the definitions of j and i, and monotonicity of f (using the fact that $\perp_D \sqsubseteq_D x$ for any x).

We leave to the reader the verification that i and j are continuous. □

Proposition 13.4.10 *F is monotone. Precisely, for every D and D',*

$$D \trianglelefteq D' \;\Rightarrow\; [D \to D] \trianglelefteq [D' \to D'].$$

Proof. By hypothesis, we have two continuous functions

$$i : D \to D' \qquad \text{and} \qquad j : D' \to D$$

such that, for all $x \in D$ and $y \in D'$,

$$j(i(x)) = x \qquad \text{and} \qquad i(j(y)) \sqsubseteq_{D'} y.$$

We prove that the functions

$$I : [D \to D] \to [D' \to D'] \qquad \text{and} \qquad J : [D' \to D'] \to [D \to D]$$

naturally defined, for all $f \in [D \to D]$ and $g \in [D' \to D']$, by

$$I(f) = i \circ f \circ j \qquad \text{and} \qquad J(g) = j \circ g \circ i$$

are continuous, and

$$J(I(f)) = f \qquad \text{and} \qquad I(J(g)) \sqsubseteq_{[D' \to D']} g.$$

Given a function $f \in [D \to D]$,

$$\begin{array}{rcl} J(I(f)) & = & J(i \circ f \circ j) \\ & = & (j \circ i) \circ f \circ (j \circ i) \\ & = & f \end{array}$$

by the definitions of I and J, and the fact that $j(i(x)) = x$ for every $x \in D$, i.e. $j \circ i = id_D$.

Given a function $g \in [D' \to D']$,

$$\begin{array}{rcl} I(J(g)) & = & I(j \circ g \circ i) \\ & = & (i \circ j) \circ g \circ (i \circ j) \\ & \sqsubseteq_{[D' \to D']} & g \end{array}$$

by the definitions of J, I and $\sqsubseteq_{[D' \to D']}$. Precisely, the last step holds because, for any $y \in D'$:

$$\begin{array}{rrcl} & i(j(y)) & \sqsubseteq_{D'} & y \\ \Rightarrow & g(i(j(y))) & \sqsubseteq_{D'} & g(y) \\ \Rightarrow & i(j(g(i(j(y))))) & \sqsubseteq_{D'} & i(j(g(y))) \quad \sqsubseteq_{D'} \quad g(y), \end{array}$$

by $i \circ j \sqsubseteq id_{D'}$, the continuity of g, i and j, and $i \circ j \sqsubseteq id_{D'}$ again.

Finally, I and J are continuous because composition preserves continuity. □

Exercise 13.4.11 *For every c.c.p.o.* D_1, D_2, D'_1 *and* D'_2*, if* $D_1 \trianglelefteq D'_1$ *and* $D_2 \trianglelefteq D'_2$*, then* $[D_1 \to D_2] \trianglelefteq [D'_1 \to D'_2]$.

The next crucial result is best stated if we introduce a special notation for D_∞, that explicitly exhibits the connection with the single elements of the chain

$$D_0 \trianglelefteq \cdots \trianglelefteq D_n \trianglelefteq \cdots$$

to which it refers. From now on, we will thus write

$$D_\infty = \bigsqcup_{n \in \omega} D_n.$$

Proposition 13.4.12 *F is continuous, in the sense that it preserves $\bigsqcup$. Precisely, if*

$$D_0 \trianglelefteq \cdots \trianglelefteq D_n \trianglelefteq \cdots$$

is any chain, then

$$F(\bigsqcup_{n \in \omega} D_n) = \bigsqcup_{n \in \omega} F(D_n),$$

i.e.

$$[D_\infty \to D_\infty] = \bigsqcup_{n \in \omega} [D_n \to D_n].$$

Proof. First, we notice that if

$$D_0 \trianglelefteq \cdots \trianglelefteq D_n \trianglelefteq \cdots ,$$

then

$$[D_0 \to D_0] \trianglelefteq \cdots \trianglelefteq [D_n \to D_n] \trianglelefteq \cdots$$

by monotonicity of F. To show that

$$[D_\infty \to D_\infty] = \bigsqcup_{n \in \omega} [D_n \to D_n]$$

we use the observation made after the proof of 13.4.8, i.e. we show that $[D_\infty \to D_\infty]$ is an upper bound to the chain $\{[D_n \to D_n]\}_{n \in \omega}$ satifying special properties.

By the hypothesis, we have continuous functions

$$i_n : D_n \to D_{n+1} \qquad \text{and} \qquad j_n : D_{n+1} \to D_n$$

with the appropriate properties. Then, as in the previous proposition, we have continuous functions

$$I_n : [D_n \to D_n] \to [D_{n+1} \to D_{n+1}]$$

and

$$J_n : [D_{n+1} \to D_{n+1}] \to [D_n \to D_n]$$

with the appropriate properties. Precisely, for any $f_n \in [D_n \to D_n]$ and n,

$$I_n(f_n) = i_n \circ f_n \circ j_n \qquad \text{and} \qquad J_n(f_{n+1}) = j_n \circ f_{n+1} \circ i_n.$$

We now need to define continuous functions

$$I_{n,\infty} : [D_n \to D_n] \to [D_\infty \to D_\infty]$$

and

$$J_{n,\infty} : [D_\infty \to D_\infty] \to [D_n \to D_n]$$

such that, for all $f_n \in [D_n \to D_n]$ and $f_\infty \in [D_\infty \to D_\infty]$:

1. $J_{n,\infty}(I_{n,\infty}(f_n)) = f_n$
2. $I_{n,\infty}(J_{n,\infty}(f_\infty)) \sqsubseteq_{[D_\infty \to D_\infty]} f_\infty$
3. $I_{n,\infty}(f_n) = I_{n+1,\infty}(I_n(f_n))$
4. $J_{n,\infty}(f_\infty) = J_n(J_{n+1,\infty}(f_\infty))$
5. $\bigsqcup_{n \in \omega} I_{n,\infty}(J_{n,\infty}(f_\infty)) = f_\infty$.

Since the proof of 13.4.7 already provides continuous functions

$$i_{n,\infty} : D_n \to D_\infty \qquad \text{and} \qquad j_{n,\infty} : D_n \to D_\infty$$

with the appropriate properties, the natural guess is to let

$$I_{n,\infty}(f_n) = i_{n,\infty} \circ f_n \circ j_{n,\infty} \qquad \text{and } J_{n,\infty}(f_\infty) = j_{n,\infty} \circ f_\infty \circ i_{n,\infty}.$$

We now check that the required properties do hold.

Condition 1 holds because

$$\begin{array}{rcl} J_{n,\infty}(I_{n,\infty}(f_n)) & = & J_{n,\infty}(i_{n,\infty} \circ f_n \circ j_{n,\infty}) \\ & = & (j_{n,\infty} \circ i_{n,\infty}) \circ f_n \circ (j_{n,\infty} \circ i_{n,\infty}) \\ & = & f_n \end{array}$$

by the definitions of $I_{n,\infty}$ and $J_{n,\infty}$, and the fact that (by 13.1)

$$j_{n,\infty} \circ i_{n,\infty} = id_{D_\infty}.$$

Condition 2 holds because

$$\begin{array}{rcl} I_{n,\infty}(J_{n,\infty}(f_\infty)) & = & I_{n,\infty}(j_{n,\infty} \circ f_\infty \circ i_{n,\infty}) \\ & = & (i_{n,\infty} \circ j_{n,\infty}) \circ f_\infty \circ (i_{n,\infty} \circ j_{n,\infty}) \\ & \sqsubseteq_{[D_\infty \to D_\infty]} & f_\infty \end{array}$$

by the definitions of $J_{n,\infty}$ and $I_{n,\infty}$, and the fact that (by 13.2)

$$i_{n,\infty} \circ j_{n,\infty} \sqsubseteq_{[D_\infty \to D_\infty]} id_{D_\infty}.$$

Condition 3 holds because

$$\begin{array}{rcl} I_{n+1,\infty}(I_n(f_n)) & = & I_{n+1,\infty}(i_n \circ f_n \circ j_n) \\ & = & (i_{n+1,\infty} \circ i_n) \circ f_n \circ (j_n \circ j_{n+1,\infty}) \\ & = & i_{n,\infty} \circ f_n \circ j_{n,\infty} \\ & = & I_{n,\infty}(f_n) \end{array}$$

by the definitions of I_n and $I_{n+1,\infty}$, the facts that (by 13.3)

$$i_{n,\infty} = i_{n+1,\infty} \circ i_n,$$

and (by 13.4)

$$j_{n,\infty} = j_n \circ j_{n+1,\infty},$$

and the definition of $I_{n,\infty}$.

Similarly, Condition 4 holds because

$$\begin{aligned} J_{n+1}(J_{n+1,\infty}(f_\infty)) &= J_{n+1}(j_{n+1,\infty} \circ f_\infty \circ i_{n+1,\infty}) \\ &= (j_n \circ j_{n+1,\infty}) \circ f_\infty \circ (i_{n+1,\infty} \circ i_n) \\ &= j_{n,\infty} \circ f_\infty \circ i_{n,\infty} \\ &= J_{n,\infty}(f_\infty). \end{aligned}$$

by the definitions of $J_{n+1,\infty}$ and J_{n+1}, 13.3, 13.4, and the definition of $J_{n,\infty}$.

Finally, Condition 5 holds because

$$\begin{aligned} f_\infty(x_\infty) &= f_\infty(\bigsqcup_{p\in\omega} i_{p,\infty}(j_{p,\infty}(x_\infty)) \\ &= \bigsqcup_{p\in\omega} f_\infty(i_{p,\infty}(j_{p,\infty}(x_\infty))) \\ &= \bigsqcup_{p\in\omega}(\bigsqcup_{q\in\omega} i_{q,\infty}(j_{q,\infty}(f_\infty(i_{p,\infty}(j_{p,\infty}(x_\infty)))))) \\ &= \bigsqcup_{n\in\omega} i_{n,\infty}(j_{n,\infty}(f_\infty(i_{n,\infty}(j_{n,\infty}(x_\infty))))) \\ &= (\bigsqcup_{n\in\omega} I_{n,\infty}(J_{n,\infty}(f_\infty)))(x_\infty) \end{aligned}$$

by the fact that (by 13.5)

$$\bigsqcup_{p\in\omega} i_{p,\infty}(j_{p,\infty}(x_\infty)) = x_\infty,$$

continuity of f_∞, 13.5 again, properties of $\bigsqcup$, and the definitions of $I_{n,\infty}$ and $J_{n,\infty}$. □

Notice that the proof simply amounts to the *identification of functions f_∞ on D_∞ with appropriate sequences of functions f_n on D_n.*

Precisely, a function f_∞ induces a function f_n defined by first taking an argument $x_n \in D_n$, embedding it into D_∞ by $i_{n,\infty}$, applying f_∞ to it, and going back to D_n by $j_{n,\infty}$, i.e.

$$f_n(x_n) = j_{n,\infty}(f_\infty(i_{n,\infty}(x_n))).$$

By definition of $J_{n,\infty}$,

$$f_n = J_{n,\infty}(f_\infty).$$

Similarly, a sequence of functions $\langle f_n\rangle_{n\in\omega}$ in $\bigsqcup_{n\in\omega}[D_n \to D_n]$ induces a function f_∞ defined piecewise as

$$f_\infty(\langle x_n\rangle_{n\in\omega}) = \langle f_n(x_n)\rangle_{n\in\omega} = \bigsqcup_{n\in\omega} i_{n,\infty}(f_n(x_n)), \tag{13.7}$$

where the last equality holds by 13.5. Since $x_n = j_{n,\infty}(x_\infty)$,

$$f_\infty(x_\infty) = \langle f_n(j_{n,\infty}(x_\infty))\rangle_{n\in\omega} = \bigsqcup_{n\in\omega} i_{n,\infty}(f_n(j_{n,\infty}(x_\infty)))$$

and, by definition of $I_{n,\infty}$,

$$f_\infty = \bigsqcup_{n\in\omega} I_{n,\infty}(f_n).$$

Exactly as the functions

$$i_\infty(\langle x_n\rangle_{n\in\omega}) = \bigsqcup_{n\in\omega} i_{n,\infty}(x_n) \qquad \text{and} \qquad j_\infty(x_\infty) = \langle j_{n,\infty}(x_\infty)\rangle_{n\in\omega}$$

defined an isomorphism between D_∞ and $\bigsqcup_{n\in\omega} D_n$, and for $x_\infty = \langle x_n\rangle_{n\in\omega}$ we obtained

$$\langle x_n\rangle_{n\in\omega} = \bigsqcup_{n\in\omega} i_{n,\infty}(x_n) \qquad \text{and} \qquad x_\infty = \langle j_{n,\infty}(x_\infty)\rangle_{n\in\omega},$$

we now have that the functions

$$I_\infty(\langle f_n\rangle_{n\in\omega}) = \bigsqcup_{n\in\omega} I_{n,\infty}(f_n) \qquad \text{and} \qquad J_\infty(f_\infty) = \langle J_{n,\infty}(f_\infty)\rangle_{n\in\omega}$$

define an isomorphism between $[D_\infty \to D_\infty]$ and $\bigsqcup_{n\in\omega}[D_n \to D_n]$, and from the identification $f_\infty = \langle f_n\rangle_{n\in\omega}$ we obtain

$$\langle f_n\rangle_{n\in\omega} = \bigsqcup_{n\in\omega} I_{n,\infty}(f_n) \qquad \text{and} \qquad f_\infty = \langle J_{n,\infty}(f_\infty)\rangle_{n\in\omega}.$$

In particular, by 13.7, *the application of a function in $[D_\infty \to D_\infty]$ to an element of D_∞ acts componentwise.*

Fixed points of the function space operator

We are finally in a position to prove the result we were looking for.

Theorem 13.4.13 (Scott [1969]) *For any c.c.p.o. D there exists a c.c.p.o. D_∞ such that*

$$D \trianglelefteq D_\infty \qquad \textit{and} \qquad D_\infty = [D_\infty \to D_\infty].$$

Proof. We mimick the proof of 13.4.1 for expansionary functions. Given D, we define:

$$\begin{array}{lcl} D_0 & = & D \\ D_{n+1} & = & [D_n \to D_n] \\ D_\infty & = & \bigsqcup_{n\in\omega} D_n. \end{array}$$

First, we show that D_∞ *exists*. For this it is enough to show that the set $\{D_n\}_{n\in\omega}$ is a chain. This is proved by induction:

- $D_0 \trianglelefteq D_1$ since
$$D_0 = D \trianglelefteq [D \to D] = D_1,$$
because F is expansionary (13.4.9).

- if $D_n \trianglelefteq D_{n+1}$, then
$$D_{n+1} = [D_n \to D_n] \trianglelefteq [D_{n+1} \to D_{n+1}] = D_{n+2},$$
because F is monotone (13.4.10).

In particular, there exist canonical embedding and projection functions
$$i_n : D_n \to D_{n+1} \qquad \text{and} \qquad j_n : D_{n+1} \to D_n,$$
defined by induction on n as follows, where $x_n \in D_n$ and $f_n : D_n \to D_n$ for every n:
$$\begin{array}{rcl} i_0(x_0) & = & \text{the constant function on } D_0 = D \text{ with value } x_0 \\ j_0(f_0) & = & f_0(\perp_D) \end{array}$$
(from the proof of 13.4.9), and
$$\begin{array}{rcccl} i_{n+1}(x_{n+1}) & = & I_n(x_{n+1}) & = & i_n \circ x_{n+1} \circ j_n \\ j_{n+1}(f_{n+1}) & = & J_n(f_{n+1}) & = & j_n \circ f_{n+1} \circ i_n \end{array}$$
(from the proof of 13.4.10). In other words, the first level is determined by the fact that F is expansionary, and determines the following ones by monotonicity.

It remains to show that D_∞ *is isomorphic to* $[D_\infty \to D_\infty]$. Recall that an element of $D_\infty = \bigsqcup_{n\in\omega} D_n$ is a sequence $\langle x_n\rangle_{n\in\omega}$ such that, for all n,
$$x_n \in D_n \qquad \text{and} \qquad x_n = j_n(x_{n+1}).$$
We get the isomorphism between D_∞ and $[D_\infty \to D_\infty]$ in two steps:

1. *identification of* $[D_\infty \to D_\infty]$ *with* $\bigsqcup_{n\in\omega} D_{n+1}$
Indeed,
$$[D_\infty \to D_\infty] = \bigsqcup_{n\in\omega} [D_n \to D_n] = \bigsqcup_{n\in\omega} D_{n+1}$$
by continuity of F (13.4.12) and definition of D_{n+1}.

   By the proof of 13.4.12, an element of $[D_\infty \to D_\infty]$ is then identifiable with a sequence $\langle f_n\rangle_{n\in\omega}$ such that, for all n,
$$f_n \in [D_n \to D_n] \qquad \text{and} \qquad f_n = J_n(f_{n+1}),$$

where the function

$$J_n : [D_{n+1} \to D_{n+1}] \to [D_n \to D_n]$$

is the canonical projection function between functions spaces defined in 13.4.10, i.e.

$$J_n(f_{n+1}) \quad = \quad j_n \circ f_{n+1} \circ i_n$$

for every $f_{n+1} \in [D_{n+1} \to D_{n+1}]$.

2. *identification of* $\bigsqcup_{n \in \omega} D_n$ *with* $\bigsqcup_{n \in \omega} D_{n+1}$
   We can now define, in the natural way,

$$\begin{array}{rcl} I(\langle x_n \rangle_{n\in\omega}) & = & \langle x_1, x_2, \dots \rangle \\ J(\langle f_n \rangle_{n\in\omega}) & = & \langle f_0(\bot_D), f_0, f_1, \dots \rangle, \end{array}$$

   where $x_n \in D_n$ and $f_n \in [D_{n+1} \to D_n] = D_{n+1}$. In other words, I forgets about the first component, while J extends a sequence in the canonical way.

   We notice that

$$I : (\bigsqcup_{n\in\omega} D_n) \to (\bigsqcup_{n} \in D_{n+1}) \qquad \text{and} \qquad J : (\bigsqcup_{n\in\omega} D_{n+1}) \to (\bigsqcup_{n\in\omega} D_n).$$

   This is obvious, because the projection functions j_n and J_n used in the definition of $\bigsqcup_{n\in\omega} D_n$ and $\bigsqcup_{n\in\omega} D_{n+1}$, respectively, are related by definition in the needed way, i.e.

$$J_n = j_{n+1}.$$

   Moreover, I and J are inverse functions:

$$I(J(\langle f_n \rangle_{n\in\omega})) = I(\langle f_0(\bot_D), f_0, f_1, \dots \rangle) = \langle f_n \rangle_{n\in\omega},$$

   and

$$J(I(\langle x_n \rangle_{n\in\omega})) = J(\langle x_1, x_2, \dots \rangle) = \langle x_1(\bot_D), x_1, \dots \rangle = \langle x_n \rangle_{n\in\omega},$$

   because $x_0 = j_0(x_1) = x_1(\bot_D)$.

Since I and J are obviously continuous, they witness that D_∞ and $[D_\infty \to D_\infty]$ are isomorphic, as we wanted to prove. □

Notice that *for any c.c.p.o. D, the c.c.p.o. D_∞ constructed above induces an extensional canonical model of the Typed Lambda Calculus*, consisting of the c.c.p.o.'s $\{D_\alpha\}_\alpha$ defined as follows:

- $D_\alpha = D$ for α atomic
- $D_{\alpha\to\beta} = [D_\alpha \to D_\beta]$.

Indeed, $D_\alpha \trianglelefteq D_\infty$ *for every* α:

- $D \trianglelefteq D_\infty$ by 13.4.13
- if $D_\alpha, D_\beta \trianglelefteq D_\infty$, then $[D_\alpha \to D_\beta] \trianglelefteq [D_\infty \to D_\infty] = D_\infty$ by 13.4.11.

Thus, although the construction of D_∞ only involves the types 0 (atomic) and $n+1 = n \to n$, its result D_∞ is actually closed under arbitrary types.

Figure 13.2 collects all facts about $D_\infty = [D_\infty \to D_\infty]$ proved so far.

Extensional models

After this long detour, we can finally go back to our original goal of finding extensional models for the Untyped Lambda Calculus.

Theorem 13.4.14 The $\boldsymbol{D_\infty}$ Model (Scott [1969]) *Any c.c.p.o.* D_∞ *such that* $D_\infty = [D_\infty \to D_\infty]$ *induces an extensional canonical model of the Untyped Lambda Calculus.*

Proof. The condition $D_\infty = [D_\infty \to D_\infty]$ means that there are continuous functions

$$I : D_\infty \to [D_\infty \to D_\infty] \qquad \text{and} \qquad J : [D_\infty \to D_\infty] \to D_\infty$$

such that $I \circ J$ and $J \circ I$ are both identities.

Consider the structure

$$\mathcal{D}_\infty = \langle D_\infty, I, J, [\![\;\;]\!]^{\mathcal{D}_\infty} \rangle,$$

where $[\![\;\;]\!]^{\mathcal{D}_\infty}$ is the canonical interpretation. In particular, given any environment ρ on D_∞,

$$[\![t]\!]_\rho = \begin{cases} \rho(x) & \text{if } t = x \\ I([\![u]\!]_\rho)([\![v]\!]_\rho) & \text{if } t = uv \\ J(\Lambda X.\, [\![u]\!]_{\rho[x:=X]}) & \text{if } t = \lambda x.\, u, \end{cases}$$

where $\Lambda X.\, [\![u]\!]_{\rho[x:=X]}$ denotes the function

$$a \in D_\infty \longmapsto [\![u]\!]_{\rho[x:=a]} \in D_\infty.$$

By 13.2.2, to prove that $\mathcal{D}_\infty$ is an extensional model it is enough to check that $[\![\;\;]\!]_\rho$ is well-defined, i.e. that $[\![t]\!]_\rho \in D_\infty$ for any t and ρ. The only differences with the proof of 10.4.3 are the absence of types, and the presence of I and J in the definition of $[\![\;\;]\!]_\rho$. The first presents no problem, and the second is taken care of by the fact that I and J are continuous, which makes the proof of the fact that $[\![u]\!]_{\rho[x:=X]}$ is a continuous function of X go through. $\square$

Given a c.c.p.o.'s $(D, \sqsubseteq_D)$, we consider

$$D_0 = D \qquad D_{n+1} = [D_n \to D_n] \qquad D_\infty = \bigsqcup_{n\in\omega} D_n,$$

If $x_n \in D_n$ and $f_n : D_n \to D_n$ for every n, then $D_0 \trianglelefteq D_1$ via

$$\begin{array}{rcl} i_0(x_0) & = & \text{the constant function on } D \text{ with value } x_0 \\ j_0(f_0) & = & f_0(\bot_D) \end{array}$$

and $D_{n+1} \trianglelefteq D_{n+2}$ via

$$\begin{array}{rcccl} i_{n+1}(x_{n+1}) & = & I_n(x_{n+1}) & = & i_n \circ x_{n+1} \circ j_n \\ j_{n+1}(f_{n+1}) & = & J_n(f_{n+1}) & = & j_n \circ f_{n+1} \circ i_n. \end{array}$$

Since

$$[D_\infty \to D_\infty] = \bigsqcup_{n\in\omega} D_{n+1}$$

by continuity of $F(X) = [X \to X]$,

$$\begin{array}{rcl} \langle x_n \rangle_{n\in\omega} \in D_\infty & \Leftrightarrow & (\forall n)(x_n \in D_n \ \wedge \ x_n = j_n(x_{n+1})) \\ \langle f_n \rangle_{n\in\omega} \in [D_\infty \to D_\infty] & \Leftrightarrow & (\forall n)(f_n \in D_{n+1} \ \wedge \ f_n = J_n(f_{n+1})), \end{array}$$

and D_∞ and $[D_\infty \to D_\infty]$ are isomorphic via the continuous functions I and J defined by:

$$\begin{array}{rcl} I(\langle x_n \rangle_{n\in\omega}) & = & \langle x_1, x_2, \dots \rangle \\ J(\langle f_n \rangle_{n\in\omega}) & = & \langle f_0(\bot_D), f_0, f_1, \dots \rangle. \end{array}$$

Moreover,

$$\langle f_n \rangle_{n\in\omega}(\langle x_n \rangle_{n\in\omega}) = \langle f_n(x_n) \rangle_{n\in\omega} = \bigsqcup_{n\in\omega} i_{n,\infty}(f_n(x_n)).$$

Figure 13.2: Definition and properties of $D_\infty = [D_\infty \to D_\infty]$

13.5 Retracts

So far we have constructed two kinds of models of the Untyped Lambda Calculus: a nonextensional one based on continuity on $\mathcal{P}(\omega)$, and an extensional one based on an inverse limit construction. We will now show how the inverse limit construction can be performed inside $\mathcal{P}(\omega)$ itself.

Recall that D_∞ was obtained by iterating the function

$$F(X) = [X \to X],$$

starting with any c.c.p.o. D:

$$\begin{array}{lcl} D_0 & = & D \\ D_{n+1} & = & F(D_n), \end{array}$$

and showing that the iteration process has a fixed point D_∞.

The idea now is that, since $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ is embedded in $\mathcal{P}(\omega)$ in a canonical way, for any c.c.p.o. X embedded in $\mathcal{P}(\omega)$ the c.c.p.o. $F(X) = [X \to X]$ is also embedded in $\mathcal{P}(\omega)$, and the construction of D_∞ can then be seen as taking place inside $\mathcal{P}(\omega)$. In particular, we can start the iteration of F with $\mathcal{P}(\omega)$ itself:

$$\begin{array}{lcl} D_0 & = & \mathcal{P}(\omega) \\ D_{n+1} & = & F(D_n). \end{array}$$

The limit D_∞ of this iteration will then be a c.c.p.o. embedded in $\mathcal{P}(\omega)$, and such that $D_\infty = [D_\infty \to D_\infty]$.

Plan of the proof

Our plan is strictly similar to the one implemented for the construction of D_∞. In particular, we will deal with special c.c.p.o.'s (called **retracts**) that are embedded in $\mathcal{P}(\omega)$ in a canonical way, and proceed as follows:

1. consider the collection of all retracts
2. define a partial ordering on retracts
3. show that the class of retracts with the associated partial ordering is a c.c.p.o.
4. show that F is continuous as a function on retracts.

Notice that, unlike the case of D_∞, the notions here involved (such as partial ordering, c.c.p.o., continuity) are taken literally, and not only as guiding analogies. The idea is that, because of its universal property, $\mathcal{P}(\omega)$ can be seen as a real c.c.p.o. of c.c.p.o.'s, and the method of retracts is the real fixed point construction that we could only mimic (in categorical terms) in the construction of D_∞.

Retracts of c.c.p.o.'s

Recall that, by definition 13.4.4, if $(D, \sqsubseteq_D)$ and $(R, \sqsubseteq_R)$ are two c.c.p.o's, then $R \trianglelefteq D$ if and only if there are two continuous functions

$$i : R \longrightarrow D \qquad \text{and} \qquad j : D \longrightarrow R$$

such that, for all $x \in R$ and $y \in D$,

$$j(i(x)) = x \qquad \text{and} \qquad i(j(y)) \sqsubseteq_D y.$$

A particularly nice case is obtained when $R \subseteq D$, i is the *identity function*, i.e. $i(x) = x$ for all $x \in R$, and $\sqsubseteq_R$ is the restriction of $\sqsubseteq_D$ to R. In this case we can drop the subscript of $\sqsubseteq_R$ and $\sqsubseteq_D$, since the two coincide when they are both defined. Moreover, the following property holds:

- *j is idempotent, i.e. $j(j(y)) = j(y)$ for all $y \in D$*
  Since $j \circ i$ is the identity, j is onto R. Let $x \in R$. Since $i(x) = x$, then $j(i(x)) = j(x)$. Since $j(i(x)) = x$, then $j(x) = x$. Thus j is the identity on its range. But a typical element of the range of j has the form $j(y)$ for $y \in D$, and thus $j(j(y)) = j(y)$.

The next definition expresses the conditions above in terms of D and j alone.

Definition 13.5.1 *Given a c.c.p.o. $(D, \sqsubseteq)$ and a continuous function*

$$j : D \longrightarrow D,$$

we say that j is a **retraction** *of D if j is idempotent, and then its range R_j is called a* **retract** *of D.*

If a retraction j satisfies the additional condition $j(y) \sqsubseteq y$, it is called a **projection**. If it satisfies the dual condition $y \sqsubseteq j(y)$, it is called a **closure**. The particular retraction used in the proof of 13.5.12 will be a closure.

The next observation characterizes retracts in terms of retractions.

Proposition 13.5.2 *If j is a retraction of D, the associated retract R_j is the set of fixed points of j.*

Proof. If $y \in D$ is a fixed point of j, then $j(y) = y$. Thus y belongs to the range of j.

Conversely, if y is in the range of j, then $y = j(x)$ for some $x \in D$, and

$$j(y) = j(j(x)) = j(x) = y$$

because j is idempotent. Thus y is a fixed point of j. □

Each retract can be seen as a c.c.p.o. in a canonical way.

Corollary 13.5.3 *A retract R_j of D is a c.c.p.o. w.r.t. the partial ordering induced by $\sqsubseteq_D$.*

Proof. If $x_0 \sqsubseteq x_1 \sqsubseteq \cdots$ is a chain of elements of R_j, it is enough to show that $\bigsqcup_{n\in\omega} x_n$, which is in D, already belongs to R_j. By the previous proposition, it is enough to show that it is a fixed point of j. And

$$j(\bigsqcup_{n\in\omega} x_n) = \bigsqcup_{n\in\omega} j(x_n) = \bigsqcup_{n\in\omega} x_n$$

because j is continuous, and $j(x_n) = x_n$ (since $x_n \in R_j$ by hypothesis, and hence it is a fixed point of j). □

Notice that the least elements of D and R_j are not necessarily the same. Actually, $\bot_{R_j}$ is the least fixed point of j.

The last two results, showing that every retraction can be seen as a c.c.p.o. in a canonical way, throught its range, allow us to restrict attention to retractions.

Definition 13.5.4 $\boldsymbol{\mathcal{R}_D}$ *is the set of all retractions of D.*

The next result shows that the collection $\mathcal{R}_D$ is itself a c.c.p.o. The restriction to $\mathcal{R}_D$ will then allow us to avoid the technical work involved in 13.4.7 and 13.4.8, that was needed to show that the collection of *all* c.c.p.o.'s was a 'c.c.p.o.' itself.

Proposition 13.5.5 *$\mathcal{R}_D$ is a c.c.p.o. w.r.t. the partial ordering induced by $\sqsubseteq_{[D\to D]}$.*

Proof. Since j is a retraction if and only if it is idempotent, i.e. $j \circ j = j$, retractions are exactly the fixed points of the selfcomposition function

$$S : [D \to D] \to [D \to D]$$

defined by

$$S(f) = f \circ f.$$

S is obviously continuous (for a similar proof, see 13.5.7). As in the previous proposition, the set of fixed points of a continuous function is still a c.c.p.o. w.r.t. the induced partial ordering. Then $\mathcal{R}_D$ is a c.c.p.o. w.r.t. the pointwise partial ordering of $[D \to D]$. □

In particular, by 13.4.1, any continuous function on $\mathcal{R}_D$ has a least fixed point, and every expansionary continuous function on $\mathcal{R}_D$ has a fixed point above any of its elements.

Function spaces of retracts

The next result is a version of 13.4.10 for retracts, and shows by a similar proof that if R_j is a retract of D, then the function space $[R_j \to R_j]$ is a retract of $[D \to D]$ in a natural way. As usual, we work with retractions.

Proposition 13.5.6 *If j is a retraction of D, then the function $j \to j$ defined by*

$$(j \to j)(f) = j \circ f \circ j$$

for every $f \in [D \to D]$, is a retraction of $[D \to D]$, and $R_{j \to j} = [R_j \to R_j]$.

Proof. For simplicity of notations, in this proof we write J for $j \to j$ and R_J for $R_{j \to j}$.

J is obviously continuous, and it is a retraction of $[D \to D]$. Indeed, for every $f \in [D \to D]$,

$$J(J(f)) = j \circ (j \circ f \circ j) \circ j = (j \circ j) \circ f \circ (j \circ j) = j \circ f \circ j = J(f)$$

because j is a retraction of D, and hence idempotent.

There are natural maps

$$F : R_J \to [R_j \to R_j] \qquad \text{and} \qquad G : [R_j \to R_j] \to R_J,$$

defined as follows.

- By definition, a typical element $J(f)$ of R_J is a function from D to R_j, and it is defined on all $x \in D$. $J(f)$ can be naturally seen as a function from R_j to R_j, by restricting it to R_j. Since $j(x) = x$ if $x \in R_j$ (by 13.5.2), this corresponds to considering the function defined as follows, for every $x \in R_j$:

$$F(J(f)) = j \circ f.$$

- Conversely, a function $g \in [R_j \to R_j]$ can be turned into an element of R_J by extending it into a function from D to D by using j first, and then applying J to it:

$$G(g) = J(g \circ j).$$

The two functions F and G are obviously continuous, and they are the inverse one of the other:

- if $f \in [D \to D]$, then

$$G(F(J(f))) = G(j \circ f) = J(j \circ f \circ j) = J(J(f)) = J(f)$$

because J is a retraction;

- if $g \in [R_j \to R_j]$, then

$$F(G(g)) = F(J(g \circ j)) = j \circ g \circ j.$$

This coincides with g on R_j, since for $x \in R_j$,

$$j(g(j(x)) = j(g(x)) = g(x)$$

because j is the identity on R_j, and both x and $g(x)$ are in R_j. □

The next result is a version of 13.4.12.

Proposition 13.5.7 *For any c.c.p.o. D, the function*

$$F : \mathcal{R}_D \to \mathcal{R}_{[D\to D]}$$

defined, in the notation of the previous proposition, by

$$F(j) = j \to j,$$

is continuous.

Proof. Recall that, by 13.5.5, the partial order of $\mathcal{R}_D$ is induced by the pointwise ordering of $[D \to D]$. We need to prove the following:

1. *F is monotone*
Let $j_1 \sqsubseteq_{[D\to D]} j_2$ and $f \in [D \to D]$. Then, for every $x \in D$,

$$\begin{array}{rrcl} & j_1(x) & \sqsubseteq_D & j_2(x) \\ \Longrightarrow & f(j_1(x)) & \sqsubseteq_D & f(j_2(x)) \\ \Longrightarrow & j_1(f(j_1(x))) & \sqsubseteq_D & j_2(f(j_2(x))) \\ \Longrightarrow & (j_1 \to j_1)(f)(x) & \sqsubseteq_D & (j_2 \to j_2)(f)(x) \end{array}$$

by the hypothesis on j_1 and j_2, and monotonicity of f and j_1. Then, for every $f \in [D \to D]$,

$$(j_1 \to j_1)(f) \sqsubseteq_{[D\to D]} (j_2 \to j_2)(f),$$

i.e.

$$F(j_1) = (j_1 \to j_1) \sqsubseteq_{[\,[D\to D]\to[D\to D]\,]} (j_2 \to j_2) = F(j_2).$$

2. *F preserves l.u.b.'s of chains*
Let

$$j_0 \sqsubseteq_{[D\to D]} j_1 \sqsubseteq_{[D\to D]} \cdots$$

be a chain of retracts in $\mathcal{R}_D$. By definition,

$$F(j)(f) = (j \to j)(f) = j \circ f \circ j.$$

Then, for every $f \in [D \to D]$ and $x \in D$,

$$\begin{aligned}
F(\bigsqcup_{n\in\omega}^{[D\to D]} j_n)(f)(x) &= ((\bigsqcup_{n\in\omega}^{[D\to D]} j_n) \circ f \circ (\bigsqcup_{n\in\omega}^{[D\to D]} j_n))(x) \\
&= (\bigsqcup_{n\in\omega}^{[D\to D]} j_n)(f((\bigsqcup_{n\in\omega}^{[D\to D]} j_n)(x))) \\
&= (\bigsqcup_{n\in\omega}^{[D\to D]} j_n)(f(\bigsqcup_{q\in\omega}^{[D\to D]} j_q(x))) \\
&= (\bigsqcup_{n\in\omega}^{[D\to D]} j_n)(\bigsqcup_{q\in\omega}^{[D\to D]} f(j_q(x))) \\
&= \bigsqcup_{p\in\omega}^{[D\to D]} (\bigsqcup_{q\in\omega}^{[D\to D]} j_p(f(j_q(x)))) \\
&= \bigsqcup_{n\in\omega}^{[D\to D]} j_n(f(j_n(x))) \\
&= \bigsqcup_{n\in\omega}^{[D\to D]} F(j_n)(f)(x)
\end{aligned}$$

by the definitions of F, $\circ$ and $\bigsqcup_{n\in\omega}^{[D\to D]}$, continuity of f, definition of $\bigsqcup_{n\in\omega}^{[D\to D]}$, continuity of j_p, monotonicity of j_q, f and j_p (which implies

$$j_p(f(j_q(x))) \sqsubseteq_D j_n(f(j_n(x)))$$

for any $n \geq p, q$), and definition of F. $\Box$

Exercises 13.5.8 a) *If j_1 and j_2 are retractions of D_1 and D_2, then the function $j_1 \to j_2$ defined by $(j_1 \to j_2)(f) = j_2 \circ f \circ j_1$, for every $f \in [D_1 \to D_2]$, is a retraction of $[D_1 \to D_2]$, and $R_{j_1 \to j_2} = [R_{j_1} \to R_{j_2}]$.*

b) *If $j_1 \sqsubseteq j_1'$ and $j_2 \to j_2'$, then $j_1 \to j_2 \sqsubseteq j_1' \to j_2'$.*

Retracts of reflexive c.c.p.o.'s

The next observation will allow us to connect retracts of D and $[D \to D]$.

Proposition 13.5.9 *For any c.c.p.o. D, the retract of a retract of D is still a retract of D. In other words, $\mathcal{R}_D$ is closed under retracts.*

Proof. Let $j_1 : D \to D$ be a retract of D, with $R_1 = j_1(D)$, and $j_2 : R_1 \to R_1$ be a retract of R_1, with $R_2 = j_2(R_1)$. Then $R_2 = j_2(j_1(D))$, and $j_2 \circ j_1 : D \to D$ is a continuous function, being the composition of two continuous functions.

It is thus enough to show that $j_2 \circ j_1$ is still a retraction, since R_2 is the image of D under it. Indeed,

$$\begin{array}{rcl} ((j_2 \circ j_1) \circ (j_2 \circ j_1))(x) & = & (j_2(j_1(j_2(j_1(x))))) \\ & = & (j_2(j_2(j_1(x)))) \\ & = & j_2(j_1(x)) \\ & = & (j_2 \circ j_1)(x) \end{array}$$

because j_1 is the identity on R_1, and hence on $j_2(j_1(x))$, and j_2 is idempotent. □

Since we have seen that function spaces of retracts of D are retracts of $[D \to D]$, and that retracts of retracts of D are retracts of D, $\mathcal{R}_D$ is closed under function spaces, whenever $[D \to D]$ is itself a retract of D.

Definition 13.5.10 *A c.c.p.o. D is* **reflexive** *if $[D \to D]$ is a retract of D.*

The next result provides a weak analogue of 13.4.13, but sufficient for our purposes.

Proposition 13.5.11 *For any reflexive c.c.p.o. D, if j is a retraction such that $j \sqsubseteq_{[D\to D]} (j \to j)$, then there exists a retraction J such that*

$$j \sqsubseteq_{[D\to D]} J \qquad \textit{and} \qquad R_J = [R_J \to R_J].$$

Proof. By 13.5.7, the function

$$F(j) = j \to j$$

is continuous as a function from $\mathcal{R}_D$ to $\mathcal{R}_{[D\to D]}$. Since D is reflexive, F can actually be seen as a continuous function from $\mathcal{R}_D$ to itself, since $j \to j$ (a retract of $[D \to D]$, hence of a retract of D) can be seen as a retract of D. We can then apply directly (the proof of) 13.4.1. □

Retracts of the Graph Model

It now remains to apply the general result just proved, by finding a reflexive c.c.p.o. D and a retraction $j \sqsubseteq_{[D\to D]} (j \to j)$ on it such that the fixed point J of F above j is not trivial.

Proposition 13.5.12 *$\mathcal{P}(\omega)$ is a reflexive c.c.p.o.*

Proof. To show that $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ is a retract of $\mathcal{P}(\omega)$, consider the functions

$$I : \mathcal{P}(\omega) \to [\mathcal{P}(\omega) \to \mathcal{P}(\omega)] \qquad \text{and} \qquad J : [\mathcal{P}(\omega) \to \mathcal{P}(\omega)] \to \mathcal{P}(\omega)$$

defined in 13.3.2 as follows:

$$J(f) = \bigoplus_{n \in \omega} f(u_n) \qquad \text{and} \qquad I(A) = \Lambda X. \bigcup_{u_n \subseteq X} (A)_n.$$

We let

$$j(A) = J(I(A)),$$

i.e. we associate with every set A the graph of the continuous function coded by A (recall that in general we only have $J(I(A)) \supseteq A$). Then j is continuous, being the composition of two continuous functions, and idempotent. Indeed,

$$j(j(A)) = J(I(J(I(A)))) = J(I(A)) = j(A)$$

because, as proved in 13.3.2, $I \circ J$ is the identity. □

Exercise 13.5.13 *$\mathcal{E}$ is a reflexive c.c.p.o.* (Hint: associate with every r.e. set A the graph of the effective continuous function coded by A.)

We have thus identified $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ with the subset of $\mathcal{P}(\omega)$ consisting of all graphs of continuous functions, i.e. all sets A such that $A = J(I(A))$. Notice the following properties:

1. *the orders of $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ as a function space and as a retract coincide*
   Recall that the order of $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ as a function space is pointwise inclusion of the values, while its order as a retract is the inclusion induced from $\mathcal{P}(\omega)$. Let A and B be graphs of continuous functions, i.e.

$$A = J(I(A)) = \bigoplus_{n \in \omega} I(A)(u_n) \text{ and } B = J(I(B)) = \bigoplus_{n \in \omega} I(B)(u_n).$$

   Then

$$\begin{aligned} A \subseteq B &\Rightarrow (\forall m)((A)_m \subseteq (B)_m) \\ &\Rightarrow (\forall n)(I(A)(u_n) \subseteq I(B)(u_n)) \\ &\Rightarrow (\forall X)(I(A)(X) \subseteq I(B)(X)) \end{aligned}$$

   by definition of I. Conversely,

$$\begin{aligned} (\forall X)(I(A)(X) \subseteq I(B)(X)) &\Rightarrow (\forall n)(I(A)(u_n) \subseteq I(B)(u_n)) \\ &\Rightarrow \bigoplus_{n \in \omega} I(A)(u_n) \subseteq \bigoplus_{n \in \omega} I(B)(u_n) \\ &\Rightarrow J(I(A)) \subseteq J(I(B)) \\ &\Rightarrow A \subseteq B \end{aligned}$$

by definition of J.

2. *a retract of* $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ *can be identified with a retract of* $\mathcal{P}(\omega)$
Since $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ is a retract of $\mathcal{P}(\omega)$, we can compose the retractions and identify a set of continuous functions with the set of their graphs.

3. *a retract of* $\mathcal{P}(\omega)$ *can be identified with an element of* $\mathcal{P}(\omega)$
We can code a retract by the retraction defining it, as a continuous function on $\mathcal{P}(\omega)$.

4. $\mathcal{P}(\omega)$ *is smaller than* $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ *in the order of retractions*
$\mathcal{P}(\omega)$ is defined by the identity function $\Lambda X.X$, and $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ by the function $\Lambda X.J(I(X))$. Now $J(I(X)) \supseteq X$ for every X, and thus $\Lambda X.\,X$ is smaller than $\Lambda X.\,J(I(X))$ in the pointwise ordering of functions, which induces the ordering of retractions.

We can now prove a weak analogue of 13.4.13.

Theorem 13.5.14 *There is a nontrivial retract* R *of* $\mathcal{P}(\omega)$ *such that* $R = [R \to R]$.

Proof. It is enough to find a nontrivial fixed point of $F(j) = j \to j$, i.e. a nontrivial retraction J such that $J = J \to J$. Then, by 13.5.6,

$$R_J = R_{J \to J} = [R_J \to R_J].$$

As a first trial, we can start from the retraction

$$j = \Lambda X.\,\bot_{\mathcal{P}(\omega)} = \Lambda X.\,\emptyset,$$

whose associated retract is $\{\emptyset\}$. This is the least element of $\mathcal{R}_{\mathcal{P}(\omega)}$, and thus its bottom. As usual in similar situations, the least fixed point of F above the bottom is the bottom itself, and we thus get a trivial fixed point.

As a second trial, we can consider the retraction

$$j = \Lambda X.\,X,$$

whose associated retract is $\mathcal{P}(\omega)$. Since, as we noted above, $\mathcal{P}(\omega)$ is smaller than $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ as a retract, we have

$$j \sqsubseteq_{[\mathcal{P}(\omega)\to\mathcal{P}(\omega)]} j \to j,$$

and 13.5.11 ensures that there is a fixed point J of F above j.

To argue that such a fixed point J is nontrivial we have to look more closely at the construction that produces it, as the l.u.b. in the pointwise order of retracts of the chain

$$j_0 = j \qquad \text{and} \qquad j_{n+1} = F(j_n) = j_n \to j_n.$$

It is enough to show that the associated retract has at least two elements. For example, $\emptyset$ and ω are such elements, since

$$J(I(\emptyset)) = \emptyset \qquad \text{and} \qquad J(I(\omega)) = \omega.$$

They are the graphs of the functions $\Lambda X.\,\bot$ and $\Lambda X.\,\top$ at every level of the construction. □

Intuitively, the construction starts with $\mathcal{P}(\omega)$, and at each level it considers only the graphs of continuous functions on the previous level. In particular, the first level gives the graphs of the continuous functions on $\mathcal{P}(\omega)$, the second level the graphs of the continuous functions that map graphs to graphs (at the first level), and so on.

From the general result 13.4.14 we can now deduce that *the retract R induces an extensional canonical model of the Untyped Lambda Calculus*.

Moreover, *R also induces an extensional canonical model of the Typed Lambda Calculus*, consisting of the retracts $\{R_\alpha\}_\alpha$ of $\mathcal{P}(\omega)$ defined as follows:

- $R_\alpha = \mathcal{P}(\omega)$ for α atomic
- $R_{\alpha\to\beta} = [R_\alpha \to R_\beta]$.

Notice that R_α *is smaller than R in the order of retractions, for every α*. Indeed:

- $\mathcal{P}(\omega)$ is smaller than R by the proof of 13.5.14
- if R_α, R_β are smaller than R, then $[R_\alpha \to R_\beta]$ is smaller than $[R \to R] = R$ by 13.5.8.b.

Sierpinski Spaces ⋆

The crucial results 13.3.2 and 13.5.12 are based on the fact that $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ is naturally embedded in $\mathcal{P}(\omega)$. This can also be obtained as a special case of a classical topological result, which we now sketch.

Consider the two-element space $2 = \{0, 1\}$ with the *Sierpinski topology*, whose open sets are $\emptyset$, $\{1\}$ and $\{0, 1\}$. This obviously defines a T_0 space (see 5.3.9), because $\{1\}$ is an open set containing 1 but not 0. Notice that the space is not T_2 (see 18.3.5), because 0 and 1 cannot be separated by disjoint open sets.

The *Sierpinski space* 2^ω is the product of the Sierpinski topology on 2. The result is homeomorphic to $\mathcal{P}(\omega)$ with the Scott topology, whose basic open sets are the sets $\{X : X \supseteq u\}$, where u is a finite set (see 6.3.9.c). Again, this obviously defines a T_0 space, because if two sets A and B are distinct, then they differ on some element x, and the basic open set $\{X : x \in X\}$ contains one of A and B, but not the other. Again, the space is not T_2, because the sets $\emptyset$ and ω cannot be separated by disjoint open sets.[6]

The next result shows that $\mathcal{P}(\omega)$ is not only a T_0 space, but a universal one.

Proposition 13.5.15 *Every T_0 space with a countable basis of open sets is homeomorphic to a subspace of $\mathcal{P}(\omega)$.*

Proof. If $\mathcal{T}$ has a countable basis, it is possible to enumerate the basic open sets by using natural numbers as indices. Then the function

$$J(x) = \{n : x \text{ belongs to the } n\text{-th basic open set}\}$$

maps elements of $\mathcal{T}$ to a set of natural numbers, i.e. to an element of $\mathcal{P}(\omega)$. And if $\mathcal{T}$ is T_0, then J is one-one.

Given the basic open set $U = \{X : X \supseteq u\}$ of $\mathcal{P}(\omega)$, then $J^{-1}(U)$ is the intersection of the finitely many basic open sets whose indices are in u. Since a topology is closed under finite intersections, $J^{-1}(U)$ is an open set of $\mathcal{T}$. Thus J is continuous.

Given the n-th basic open set N of $\mathcal{T}$, then

$$J(N) = \{X : n \in X\} = \{X : X \supseteq \{n\}\},$$

which is an open set of $\mathcal{P}(\omega)$. Thus J^{-1} is continuous, and $\mathcal{T}$ is homeomorphic to its image in $\mathcal{P}(\omega)$. □

More generally, the same proof shows that every T_0 space is homeomorphic to a subspace of the Sierpinski space 2^α, where α is the cardinality of the basis of the given space.

Proposition 13.5.16 *$[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ is a T_0 space with a countable basis of open sets.*

Proof. By 6.3.9.c, the basic open sets of $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ are the sets $\{f : x \in f(u)\}$, where x is a natural number and u is a finite set.

The space is obviously T_0, because if two continuous functions g and h are distinct, then they differ on some finite set u, i.e. there is some x which is in one of $g(u)$ and $h(u)$, but not in the other. Then the basic open set $\{f : x \in f(u)\}$ contains one of g and h, but not the other.

The space has a countable basis, because the basic open sets $\{f : x \in f(u_n)\}$ can be enumerated by the numbers $\langle x, n\rangle$. □

[6] The Sierpinski space 2^ω should not be confused with the *Cantor space* 2^ω, which is the product of the *discrete topology* on $2 = \{0, 1\}$, whose open sets are $\emptyset$, $\{0\}$, $\{1\}$ and $\{0, 1\}$. The Cantor space is T_2, because if two functions g and h are distinct, then they differ on some argument x, and the basic open sets $\{f : f(x) = 0\}$ and $\{f : f(x) = 1\}$ separate g and h.

If follows from 13.5.15 that $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ is homeomorphic to a subspace of $\mathcal{P}(\omega)$. Moreover, the homeomorphism provided by the proof above is nothing else than the function J used in 13.3.2, as its name implies.

Exercises 13.5.17 A universal algebraic lattice. The notion of an algebraic lattice is defined in 18.5.1. $\mathcal{P}(\omega)$ is obviously algebraic, with the finite sets as the compact elements. The next results show that $\mathcal{P}(\omega)$ is a universal algebraic lattice.

a) *Every algebraic lattice with a countable basis of compact elements is isomorphic to a sublattice of* $\mathcal{P}(\omega)$. (Hint: enumerate the compact elements by using natural numbers as indices, and define

$$J(x) = \{n : x \text{ is above the } n\text{-th compact element}\}.)$$

b) $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ *is an algebraic lattice with a countable basis of compact elements.* (Hint: the compact elements are the continuous functions that take finite values for finitely many arguments, and value $\emptyset$ for the remaining ones.)

13.6 The Least Fixed Point Operator ⋆

We now compute the interpretation of the fixed-point combinator

$$\mathcal{Y} = \lambda y.\, (\lambda x.\, y(xx))(\lambda x.\, y(xx)).$$

Since $\mathcal{Y}$ begins with a λ, intuitively its interpretation defines a function. We then actually compute $I(\llbracket \mathcal{Y} \rrbracket)$, thus characterizing the behavior of $\llbracket \mathcal{Y} \rrbracket$ as a function on $\mathcal{P}(\omega)$ and on D_∞.

The Graph Model

Recall that, in the notations of 12.2.2,

$$\mathcal{Y} = \lambda y.\, \Delta\Delta_y,$$

where $\Delta = \lambda x.\, xx$ and $\Delta_y = \lambda x.\, y(xx)$. We first notice that, by repeated applications of the definition of $\llbracket\ \rrbracket$,

$$\llbracket \Delta \rrbracket = \llbracket \lambda x. xx \rrbracket = J(\Lambda X.\, I(X)(X))$$

and

$$\llbracket \Delta_y \rrbracket = \llbracket \lambda x.\, y(xx) \rrbracket = J(\Lambda X.\, I(Y)(I(X)(X))).$$

Then, using twice the fact that $I \circ J$ is the identity,

$$I(\llbracket \mathcal{Y} \rrbracket) = (\Lambda Y.\, (\Lambda X.\, I(X)(X))(J(\Lambda X.\, I(Y)(I(X)(X)))).$$

Finally, for any $f \in [\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$,

$$\begin{aligned} I(\llbracket\mathcal{Y}\rrbracket)(J(f)) &= (\Lambda X.\, I(X)(X))(J(\Lambda X.\, I(J(f))(I(X)(X))) \\ &= (\Lambda X.\, I(X)(X))(J(\Lambda X.\, f(I(X)(X))) \\ &= G(J(F)), \end{aligned}$$

where

$$G = \Lambda X.\, I(X)(X) \qquad \text{and} \qquad F = \Lambda X.\, f(I(X)(X)).$$

We are now ready to characterize $I(\llbracket\mathcal{Y}\rrbracket)$. Notice that $[\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$ is a c.c.p.o., and thus the Fixed Point Theorem 13.4.1 holds for it. In other words, every function f on $\mathcal{P}(\omega)$ has a least fixed point. Then there is a function

$$\mathbf{Fix}_{\mathcal{P}(\omega)} : [\mathcal{P}(\omega) \to \mathcal{P}(\omega)] \to \mathcal{P}(\omega)$$

defined by:

$$\mathbf{Fix}_{\mathcal{P}(\omega)}(f) = \text{the least fixed point of } f.$$

We now prove that $I(\llbracket\mathcal{Y}\rrbracket) = \mathbf{Fix}_{\mathcal{P}(\omega)}$, thus showing in particular that $\mathbf{Fix}_{\mathcal{P}(\omega)}$ *is a continuous function*, a fact that could also be easily proved directly.

Proposition 13.6.1 (Scott [1975]) *In $\mathcal{P}(\omega)$, $I(\llbracket\mathcal{Y}\rrbracket)$ is the least fixed point operator $\mathbf{Fix}_{\mathcal{P}(\omega)}$.*

Proof. We first prove that $I(\llbracket\mathcal{Y}\rrbracket)$ is a *fixed point operator*, i.e. $I(\llbracket\mathcal{Y}\rrbracket)(J(f))$ is a fixed point of f. Since

$$I(\llbracket\mathcal{Y}\rrbracket)(J(f)) = G(J(F)),$$

it is enough to show that

$$G(J(F)) = f(G(J(F))).$$

But

$$G(J(F)) = I(J(F))(J(F)) = F(J(F))$$

by definition of G and the fact that $I \circ J$ is the identity. Thus it is enough to show that

$$F(J(F)) = f(F(J(F))).$$

And indeed,

$$F(J(F)) = f(I(J(F))J(F)) = f(F(J(F)))$$

by definition of F and the fact that $I \circ J$ is the identity. In one word, $I(\llbracket\mathcal{Y}\rrbracket)$ is a fixed point operator because $\mathcal{P}(\omega)$ is a model of the Untyped Lambda Calculus, and $\mathcal{Y}$ is a fixed point combinator.

We now prove that $I([\![\mathcal{Y}]\!])$ is the *least* fixed point operator. The idea is to show that, for any $f \in [\mathcal{P}(\omega) \to \mathcal{P}(\omega)]$,

$$f(A) = A \ \Rightarrow\ I([\![\mathcal{Y}]\!])(J(f)) \subseteq A,$$

i.e. that $I([\![\mathcal{Y}]\!])(J(f))$ is contained in every fixed point of f. Equivalently, we can show that

$$f(A) = A \ \Rightarrow\ G(J(F)) \subseteq A.$$

Let $x \in G(J(F))$. Since $G(X) = I(X)(X)$ and G is continuous, there is a finite set $u \subseteq J(F)$ such that $x \in I(u)(u)$. It is then enough to show, by induction on the code of u, that

$$u \subseteq J(F) \ \Rightarrow\ I(u)(u) \subseteq A,$$

since then $x \in A$. By definition of I,

$$I(u)(u) = \bigcup_{u_n \subseteq u} (u)_n.$$

Since $x \in I(u)(u)$, there is an n such that $u_n \subseteq u$ and $x \in (u)_n$, i.e. $\langle x, n\rangle \in u$. Then n is smaller than the code of u, and by the induction hypothesis

$$u_n \subseteq u \subseteq A \ \Rightarrow\ I(u_n)(u_n) \subseteq A.$$

But $\langle x, n\rangle \in u \subseteq J(F)$, i.e.

$$\langle x, n\rangle \in J(F) = \bigoplus_{m \in \omega} F(u_m)$$

by definition of J. Then

$$x \in F(u_n) = f(I(u_n)(u_n)) \subseteq f(A) = A,$$

by definition of F, monotonicity of f, and the fact that $I(u_n)(u_n) \subseteq A$. □

Corollary 13.6.2 *In $\mathcal{P}(\omega)$, $[\![\Delta\Delta]\!] = \emptyset$.*

Proof. Since $\Delta\Delta$ is the result of applying $\mathcal{Y}$ to the identity function, $[\![\Delta\Delta]\!]$ must be the least fixed point of the identity. But every point is a fixed point of the identity, and thus the least fixed point is the least element. □

Exercises 13.6.3 a) *The least fixed point of an effective continuous function is r.e.* (Rogers [1959]) (Hint: the proof of 13.4.1 is effective.)

b) *In $\mathcal{E}$, $I([\![\mathcal{Y}]\!])$ is the least fixed point operator* (Scott [1975]). (Hint: by 13.6.1 and part a).)

The previous results crucially depend on the choice of the coding function $\langle\ \rangle$, and may fail with a different choice.

Proposition 13.6.4 (Baeten and Boerboom [1979]) *By changing the coding function $\langle\ \rangle$ we can get a model in which $[\![\Delta\Delta]\!] \neq \emptyset$, and thus $I([\![\mathcal{Y}]\!])$ is not the least fixed point operator.*

Proof. The main observation is that

$$\{\langle x, n\rangle\} = u_n \ \Rightarrow\ x \in [\![\Delta\Delta]\!].$$

Indeed,

$$\langle x, n\rangle \in u_n \ \Rightarrow\ x \in (u_n)_n \ \Rightarrow\ x \in \bigcup_{u_m \subseteq u_n} (u_n)_m \ \Rightarrow\ x \in I(u_n)(u_n)$$

by definition of $(\ )_n$, the fact that $u_n \subseteq u_n$, and definition of I. Then

$$\langle x, n\rangle \in \bigoplus_{m\in\omega} I(u_m)(u_m) = J(\Lambda X.\, I(X)(X)) = [\![\Delta]\!].$$

It follows that $u_n = \{\langle x, n\rangle\} \subseteq [\![\Delta]\!]$ and, from $x \in (u_n)_n$,

$$x \in \bigcup_{u_n \subseteq [\![\Delta]\!]} ([\![\Delta]\!])_n = I([\![\Delta]\!])([\![\Delta]\!]) = [\![\Delta\Delta]\!].$$

The observation just made does not make use of any particular property of the coding function $\langle\ \rangle$, except for the fact that it is one-one, and it thus works for any coding function $\langle\ \rangle^*$. Since $u_{2^n} = \{n\}$, we can then ensure $[\![\Delta\Delta]\!] \supseteq \omega$, and hence $[\![\Delta\Delta]\!] = \omega$, by finding a coding function $\langle\ \rangle^*$ such that

$$\forall x \exists n (\langle x, 2^n\rangle^* = n).$$

Then for every x there is n such that $\{\langle x, 2^n\rangle^*\} = u_{2^n}$, and $x \in [\![\Delta\Delta]\!]$.

The coding function $\langle\ \rangle^*$ is obtained by inductively modifying the standard coding function $\langle\ \rangle$ in two points for each $x \in \omega$, as follows. For $x = 0$ we can let

$$\langle 0, 2^0\rangle^* = \langle 0, 1\rangle^* = 0.$$

We have now assigned the code 0 to the pair $(0, 1)$. Originally, 0 was instead assigned to $(0, 0)$. To avoid destroying one-onenness we simply assign to $(0, 0)$ the code 2, that was originally assigned to $(0, 1)$. Thus

$$\langle 0, 0\rangle^* = 2.$$

Having modified $\langle\ \rangle$ to take care of the condition above for all numbers smaller than x, we can choose an n so large that letting

$$\langle x, 2^n\rangle^* = n \qquad \text{and} \qquad \langle (n)_1, (n)_2\rangle^* = \langle x, 2^n\rangle$$

would not modify the work done for the previous values $0, \ldots, x-1$. □

Exercise 13.6.5 *For any set A there is a coding function that produces $[\![\Delta\Delta]\!] = A$.* (Baeten and Boerboom [1979]) (Hint: if $A \neq \emptyset$, ensure $\exists n(\langle x, 2^n\rangle^* = n)$ exactly for $x \in A$.)

Actually, Baeten and Boerboom [1979] have proved that *for any closed term t there is a coding function that produces* $[\![\Delta\Delta]\!] = [\![t]\!]$. It is not enough to simply let $A = [\![t]\!]$ in the previous exercise, because modifying the coding may make $[\![\Delta\Delta]\!] = A$ but change $[\![t]\!]$ at the same time. The proof requires instead a forcing argument.

The D_∞ Model

Notice that, by repeated applications of the definition of $[\![\ \ ]\!]$,

$$[\![\lambda x.\, y(xx)]\!] = J(\Lambda X.\, I(Y)(I(X)(X))).$$

Similarly, using the fact that $I \circ J$ is the identity,

$$I([\![\mathcal{Y}]\!]) = (\Lambda Y.\, (\Lambda X.\, I(Y)(I(X)(X)))(J(\Lambda X.\, I(Y)(I(X)(X)))).$$

Finally, for any $f_\infty \in [D_\infty \to D_\infty]$,

$$\begin{array}{rcl} I([\![\mathcal{Y}]\!])(J(f_\infty)) & = & (\Lambda X.\, I(J(f_\infty))(I(X)(X)))(J(\Lambda X.\, I(J(f_\infty))(I(X)(X))) \\ & = & (\Lambda X.\, f_\infty(I(X)(X)))(J(\Lambda X.\, f_\infty(I(X)(X))) \\ & = & F_\infty(J(F_\infty)), \end{array}$$

where

$$F_\infty = \Lambda X.\, f_\infty(I(X)(X)).$$

Recalling that application on D_∞ acts componentwise, we finally reduce to

$$I([\![\mathcal{Y}]\!])(J(f_\infty)) \quad = \quad \bigsqcup_{n\in\omega} i_{n,\infty}(F_n(F_{n-1})), \tag{13.8}$$

where

$$F_n = J_{n,\infty}(F_\infty) = j_{n,\infty} \circ F_\infty \circ i_{n,\infty}.$$

Moreover, for all $n \geq -1$,

$$F_n = j_{n+1}(F_{n+1}).$$

Notice that this takes care also of the case $n = -1$, corresponding to the first component of $J(F_\infty)$. By definition, the latter is obtained by projecting F_0 by j_0. Or, equivalently, by computing F_0 on $\perp_D$.

Before proceeding further we notice that, for any $x_n \in D_n$,

$$\begin{array}{rcl} F_n(x_n) & = & j_{n,\infty}(F_\infty(i_{n,\infty}(x_n))) \\ & = & j_{n,\infty}(f_\infty(I(i_{n,\infty}(x_n))(i_{n,\infty}(x_n)))) \\ & = & f_n(i_n(x_n)(x_n)). \end{array}$$

The first equality holds by definition of F_n. The second holds by definition of F_∞. The last holds by using twice the fact that application on D_∞ acts componentwise. Notice that, by definition of $i_{n,\infty}$, the n-th and $n+1$-th components of $i_{n,\infty}(x_n)$ are x_n and $i_n(x_n)$, respectively. Since I shifts all components down by one position, $i_n(x_n)$ is the n-th component of $I(i_{n,\infty}(x_n))$.

The preceeding lines are simply applications of general facts proved in the previous subsections. The new observation is that *all components* $F_n(F_{n-1})$ *in 13.8 are determined, inductively, by* $F_0(F_{-1})$. Indeed:

$$\begin{aligned} F_{n+1}(F_n) &= f_{n+1}(i_{n+1}(F_n)(F_n)) \\ &= f_{n+1}(i_n(F_n(j_n(F_n)))) \\ &= f_{n+1}(i_n(F_n(F_{n-1}))) \end{aligned}$$

by the property of F_{n+1} just proved above, the definition of $i_{n+1} = I_n$ on D_∞, and the fact that $F_{n-1} = J_{n-1}(F_n) = j_n(F_n)$. Then, by induction,

$$F_{n+1}(F_n) = f_{n+1}(i_n(\cdots(f_1(i_0(F_0(F_{-1}))))\cdots)).$$

We are now ready to characterize $I([\![\mathcal{Y}]\!])$. Notice that $[D_\infty \to D_\infty]$ is a c.c.p.o., and thus the Fixed Point Theorem 13.4.1 holds for it. In other words, every function f_∞ on D_∞ has a least fixed point. Then there is a function

$$\mathbf{Fix}_{D_\infty} : [D_\infty \to D_\infty] \to D_\infty$$

defined by:

$$\mathbf{Fix}_{D_\infty}(f_\infty) = \text{the least fixed point of } f_\infty.$$

We now prove that $I([\![\mathcal{Y}]\!]) = \mathbf{Fix}_{D_\infty}$, thus showing in particular that $\mathbf{Fix}_{D_\infty}$ *is a continuous function*, a fact that could also be easily proved directly.

Proposition 13.6.6 (Park [1970]) *In* D_∞, $I([\![\mathcal{Y}]\!])$ *is the least fixed point operator* $\mathbf{Fix}_{D_\infty}$.

Proof. We first prove that $I([\![\mathcal{Y}]\!])$ is a *fixed point operator*, i.e. $I([\![\mathcal{Y}]\!])(J(f_\infty))$ is a fixed point of f_∞. Since

$$I([\![\mathcal{Y}]\!])(J(f_\infty)) = F_\infty(J(F_\infty)),$$

it is enough to show that

$$F_\infty(J(F_\infty)) = f_\infty(F_\infty(J(F_\infty))).$$

And, indeed,

$$\begin{aligned} F_\infty(J(F_\infty)) &= f_\infty(I(J(F_\infty))(J(F_\infty))) \\ &= f_\infty(F_\infty(J(F_\infty))) \end{aligned}$$

by definition of F_∞, i.e.

$$F_\infty = \Lambda X.\, f_\infty(I(X)(X)),$$

and the fact that $I \circ J$ is the identity. In one word, $I([\![\mathcal{Y}]\!])$ is a fixed point operator because D_∞ is a model of Lambda Calculus, and $\mathcal{Y}$ is a fixed point combinator.

We now prove that $I([\![\mathcal{Y}]\!])$ is the *least* fixed point operator. The idea is to show that

$$I([\![\mathcal{Y}]\!])(J(f_\infty)) \sqsubseteq \bigsqcup_{n\in\omega} f_\infty^{(n)}(\bot_{D_\infty}).$$

By the proof of 13.4.1, the right-hand-side is the least fixed point of f_∞. By the first part of the present proof, the left-hand-side is a fixed point of f_∞. If the previous inequality holds, then the left-hand-side is the least fixed point of f_∞.

Recall that, as noted above,

$$F_{-1} = F_0(\bot_D) \qquad \text{and} \qquad F_0(x_0) = f_0(i_0(x_0)(x_0)).$$

Moreover, i_0 associates to every element of $D_0 = D$ the constant function with that element as value. Thus

$$F_0(x_0) = f_0(i_0(x_0)(x_0)) = f_0(x_0)$$

and

$$F_0(F_{-1}) = F_0(F_0(\bot_D)) = f_0(f_0(\bot_D)) = f_0^{(2)}(\bot_D).$$

Then, at the next level,

$$F_1(F_0) = f_1(i_0(F_0(F_{-1}))) = f_1(i_0(f_0(f_0(\bot_D)))).$$

Notice that, for any $x_0 \in D_0$,

$$\begin{array}{lcl} i_0(f_0(x_0)) & = & i_0(j_1(f_1(x_0))) \\ & = & i_0(j_0(f_1(i_0(x_0)))) \\ & \sqsubseteq_{D_1} & f_1(i_0(x_0)) \end{array}$$

by $f_0 = J_0(f_1) = j_1(f_1)$, definition of j_1, and the fact that $i_0 \circ j_0$ is less than the identity. In particular, for $x_0 = f_0(\bot_{D_0})$ we get

$$i_0(f_0(f_0(\bot_{D_0}))) \sqsubseteq_{D_1} f_1(f_1(i_0(\bot_{D_0}))) = f_1^{(2)}(\bot_{D_1}),$$

because i_0 is continuous, and hence it preserves $\bot$'s. Finally,

$$F_1(F_0) \sqsubseteq_{D_1} f_1^{(3)}(\bot_{D_1})$$

by monotonicity of f_1.

In a similar way we get, in general,

$$F_n(F_{n-1}) \sqsubseteq_{D_n} f_n^{(n+2)}(\bot_{D_n})$$

and thus

$$\begin{array}{lll} i_{n,\infty}(F_n(F_{n-1})) & \sqsubseteq_{D_\infty} & i_{n,\infty}(f_n^{(n+2)}(\bot_{D_n})) \\ & = & i_{n,\infty}(j_{n,\infty}(f_\infty^{(n+2)}(\bot_{D_\infty}))) \\ & \sqsubseteq_{D_\infty} & f_\infty^{(n+2)}(\bot_{D_\infty})) \end{array}$$

by monotonicity of $i_{n,\infty}$, the facts that application on D_∞ acts componentwise and $\bot_{D_\infty} = \langle \bot_{D_n} \rangle_{n\in\omega}$, and the fact that $i_{n,\infty} \circ j_{n,\infty}$ is less than the identity. Then

$$I(\llbracket \mathcal{Y} \rrbracket)(J(f_\infty)) = \bigsqcup_{n\in\omega} i_{n,\infty}(F_n(F_{n-1}))$$

is smaller than or equal to $\bigsqcup_{n\in\omega} f_\infty^{(n+2)}(\bot_{D_\infty})$, which is equal to the least fixed point of f_∞. □

Corollary 13.6.7 *In D_∞, $\llbracket \Delta\Delta \rrbracket = \bot_{D_\infty}$.*

Proof. Since $\Delta\Delta$ is the result of applying $\mathcal{Y}$ to the identity function, $\llbracket \Delta\Delta \rrbracket$ must be the least fixed point of the identity. But every point is a fixed point of the identity, and thus the least fixed point is the least element. □

The previous results crucially depend on the choice of the embedding and projection functions i_0 and j_0, and may fail with a different choice.

Proposition 13.6.8 (Park [1970]) *By changing the embedding and projection functions*

$$i_0 : D_0 \to [D_0 \to D_0] \qquad \textit{and } j_0 : [D_0 \to D_0] \to D_0$$

we can get a model in which $\llbracket \Delta\Delta \rrbracket \neq \bot_{D_\infty}$, and thus $I(\llbracket \mathcal{Y} \rrbracket)$ is not the least fixed point operator.

Proof. Notice that

$$\llbracket \Delta\Delta \rrbracket = I(\llbracket \mathcal{Y} \rrbracket)(J(f_\infty)),$$

where f_∞ is the identity function on D_∞. In this case, each f_n is the identity on D_0. Moreover, if

$$I(\llbracket \mathcal{Y} \rrbracket)(J(f_\infty)) = F_\infty(J(F_\infty)) = \langle F_n(F_{n-1}) \rangle_{n\in\omega}$$

is $\bot_{D_\infty}$, then each component $F_n(F_{n-1})$ must be $\bot_{D_n}$. To have

$$\llbracket \Delta\Delta \rrbracket \neq \bot_{D_\infty},$$

it is thus enough to define i_0 and j_0 in such a way that $F_0(F_{-1}) \neq \perp_{D_0}$.

Let D_0 be a finite nontrivial c.c.p.o.'s, choose $d \neq \perp_{D_0}$, and define

$$\begin{array}{rcl} i_0(x_0) & = & \text{the function } z_0 \longmapsto \begin{cases} x_0 & \text{if } d \sqsubseteq_D z_0 \\ \perp_D & \text{otherwise} \end{cases} \\ j_0(f_0) & = & f_0(d). \end{array}$$

Thus $i_0(x_0)$ is piecewise constant, with value x_0 above d, and value $\perp_{D_0}$ otherwise.

Both i_0 and j_0 are continuous, because they are monotone and D_0 is finite. Moreover,

$$j_0(i_0(x_0)) = x_0$$

because the left-hand-side is the value of i_0 at d. And

$$i_0(j_0(f_0)) \sqsubseteq_{D_1} f_0$$

because the left-hand-side is piecewise constant, with value $f_0(d)$ above d, and value $\perp_{D_0}$ otherwise.

We now compute $F_0(F_{-1})$, where f_0 is the identity function on D_0. First we notice that

$$F_{-1} = F_0(d) = f_0(i_0(d)(d)) = f_0(d) = d$$

by definition of F_{-1}, properties of F_0, the fact that $i_0(d)$ has value d on d, and the fact that f_0 is the identity D_0. Then

$$F_0(F_{-1}) = F_0(d) = d \neq \perp_{D_0},$$

as wanted. □ æ

Chapter 14

Cartesian Closed Monoids

æ

Chapter 15

Computability

In Chapter 11 we have seen that in the Typed Lambda Calculus we can represent exactly the piecewise polynomials. Since this is quite limited from a computational point of view, a number of stronger type systems have been introduced. However, all typed systems share the fundamental limitation of not being able to represent all the effectively computable functions. It is only in the Untyped Lambda Calculus that we reach the fullest possible computational power, which brings with it an explosion of the computational complexity.

We will provide background on the recursive functions, and refer to Odifreddi [1989] for a detailed treatment.

15.1 Representability

We have already introduced the natural numbers in the Typed Lambda Calculus, by means of numerals representing iterations. We keep the same representation in the Untyped Lambda Calculus, with the appropriate modifications allowed by the lack of types.

Definition 15.1.1 (Peano [1891], Wittgenstein [1922], Church [1933]) *The numeral $\overline{n}$ is the λ-term that produces n iterations of its first argument on the second, i.e.*

$$\overline{n} = \lambda f x.\, f^{(n)} x,$$

where

$$f^{(n)} x = \underbrace{f(\cdots(f}_{n \text{ times}}\, x)\cdots).$$

Notice that we now have only one numeral $\overline{n}$ for every number n, while in the Typed Lambda Calculus we had infinitely many $\overline{n}^{N_\alpha}$, one for each type α.

Since numerals are terms in normal form, by the Church-Rosser Theorem they are β-different if they are *syntactically* different. Thus

$$n \neq m \;\Rightarrow\; \overline{n} \neq_\beta \overline{m}.$$

We also keep the same notion of representable function introduced in the Typed Lambda Calculus, again with the appropriate simplifications allowed by the lack of types.

Definition 15.1.2 *A n-ary function f is* **representable** *in the Untyped Lambda Calculus if there is a closed term F such that, for every x_1, ..., x_n and y,*

$$f(x_1, \ldots, x_n) = y \;\Leftrightarrow\; F\overline{x}_1 \cdots \overline{x}_n =_\beta \overline{y}.$$

Examples

A simple observation immediately takes us beyond the functions representable in the Typed Lambda Calculus.

Proposition 15.1.3 (Church [1933]) *The constant functions, as well as sum, product, exponential and superexponential, are representable in the Untyped Lambda Calculus.*

Proof. By erasing the types in the proofs of Section 11.2. For example, the exponential function is representable by the term $\overline{n}\,\overline{m}$, because

$$f^{(m^n)}(x) = f^{\overbrace{m \cdots m}^{n \text{ times}}}(x) \;\;\Rightarrow\;\; \overline{m^n} f x =_\beta (\overline{n}\,\overline{m}) f x.$$

Notice that the difficulties related to the fact that in the Typed Lambda Calculus the terms $\overline{n}$ and $\overline{m}$ cannot not have the same type, obviously disappear when there are no types. □

The previous examples straightforwardly exploited the power of iteration embedded in the very definition of a numeral. The next example is subtler, and it originally constituted a stumbling block whose solution opened the way to the representability of all recursive functions.

Proposition 15.1.4 (Kleene [1935]) *The predecessor function is representable in the Untyped Lambda Calculus.*

Proof. The predecessor of n is the second component of the n-th iteration of the function t defined as follows:

$$t(\langle x, y \rangle) = \langle x + 1, x \rangle,$$

with the iteration started on the pair $\langle 0, 0\rangle$. Since iteration is given for free by the definition of numerals, and the representable functions contain the successor and are closed under composition, it is enough to show how to represent coding and decoding functions of pairs.

A coding function is obtained by analogy with the representation of numbers, as follows:

$$\overline{\langle n, m\rangle} = \lambda fgx.\, f^{(n)} g^{(m)} x.$$

Thus the coding function is represented by

$$\lambda yzfgx.\, yf(zgx).$$

Decoding is then immediate. If $\mathbf{I} = \lambda x.\, x$, then

$$\lambda fx.\, \overline{\langle n, m\rangle} f\mathbf{I}x =_\beta \lambda fx.\, f^{(n)} x = \overline{n}$$

and

$$\lambda gx.\, \overline{\langle n, m\rangle}\mathbf{I}gx =_\beta \lambda gx.\, g^{(m)} x = \overline{m}.$$

Thus the decoding functions are represented by

$$\lambda yfx.yf\mathbf{I}x \qquad \text{and} \qquad \lambda ygx.y\mathbf{I}gx. \quad \Box$$

Corollary 15.1.5 *Subtraction, as well as the characteristic functions of the ordering and equality, are representable in the Untyped Lambda Calculus.*

Proof. We have proved in 11.3.2 that:

- if the predecessor is representable, then so is subtraction
- if subtraction is representable, then so is the ordering
- if the ordering is representable, then so is equality.

In the Typed Lambda Calculus, the nonrepresentability of equality successively implied the nonrepresentability of the ordering, subtraction and predecessor. In the Untyped Lambda Calculus, the representability of predecessor successively implies the representability of subtraction, the ordering and equality. □

Recursive functions

The next definition introduces the class of recursive functions in a particularly convenient way, equivalent to the usual ones in the literature (see Odifreddi [1989], II.2.15).

Definition 15.1.6 *The class of* **recursive functions** *is the smallest class:*

- *containing the constant and projection functions, as well as successor and predecessor*
- *closed under composition, definition by cases and fixed point definitions.*

By a projection function we mean any function $\mathcal{I}_i^n$ of n variables that takes the i-th one as a value. Such functions are useful in variable manipulations, such as interchange, identification and introduction.

Proposition 15.1.7 (Church [1933], Rosser [1935], Kleene [1935], [1936]) *All recursive functions are representable in the Untyped Lambda Calculus.*

Proof. Representability of the constant functions, successor and predecessor follows from the examples above.

Representability of the projections, as well as closure under composition and definition by cases, follows from the proof of 11.2.6.

Closure under fixed point definitions follows from 12.2.2, which provides a fixed point operator. □

The Characterization Theorem

The previous result makes it very difficult to produce natural examples of nonrepresentable functions, since they would have to be nonrecursive, and hence not effectively computable by Church's Thesis (see Odifreddi [1989], Section I.8 for a discussion of it). We content ourselves with proving the converse of 15.1.7, thus characterizing the class of functions representable in the Untyped Lambda Calculus.

Proposition 15.1.8 (Church [1936], Kleene [1936]) *If a function is representable in the Untyped Lambda Calculus, then it is recursive.*

Proof. Suppose f is representable by F, i.e.

$$f(x_1, \ldots, x_n) = y \Leftrightarrow F\overline{x}_1 \cdots \overline{x}_n =_\beta \overline{y}$$

for every $x_1, \ldots, x_n$ and y. To compute $f(x_1, \ldots, x_n)$ we generate all possible β-equalities in a systematic fashion, using the rules 12.1.4, 12.1.5 and 12.1.6, until we find a numeral $\overline{y}$ such that $F\overline{x}_1 \cdots \overline{x}_n =_\beta \overline{y}$. By the hypothesis on F, the numeral $\overline{y}$ exists and is unique. Moreover, $y = f(x_1, \ldots, x_n)$. Then f is effectively computable, and hence recursive by Church's Thesis. □

Corollary 15.1.9 Characterization of the Representable Functions. *The functions representable in the Untyped Lambda Calculus are exactly the recursive functions.*

Notice that the characterization just proved confines itself to *total* functions. However, a similar result holds for *partial* functions as well. In particular, the fact that β-equality is a recursively enumerable relation implies that a partial function representable in the Untyped Lambda Calculus has a recursively enumerable graph, which is a property equivalent to being partial recursive.

Logical operators

Exactly as we mirrored numbers in the Lambda Calculus by terms, we can mirror truth-values as well. The particular choice of representatives is made with the purpose of making the proof of 15.1.12 trivial.

Definition 15.1.10 *The truth-values T and F are represented by the terms*

$$\overline{T} = \lambda xy.\, x \qquad \textit{and} \qquad \overline{F} = \lambda xy.\, y.$$

Since $\overline{T}$ and $\overline{F}$ are in normal form and syntactically different, they are β-different by the Church-Rosser Theorem.

Having represented truth-values, we can now look for representations of truth-valued functions.

Definition 15.1.11 *An n-ary truth-valued function f is* **representable** *in the Untyped Lambda Calculus if there is a closed term F such that, for all possible truth-values $x_1, \ldots, x_n$ and y,*

$$f(x_1, \ldots, x_n) = y \;\Leftrightarrow\; F\,\overline{x}_1 \cdots \overline{x}_n =_\beta \overline{y}.$$

Proposition 15.1.12 *The classical connectives are representable in the Untyped Lambda Calculus.*

Proof. We first consider the "if then else" operator δ defined (as in 11.2.4) as follows:

$$\delta(x, y, z) = \begin{cases} y & \text{if } x = T \\ z & \text{if } x = F. \end{cases}$$

By definition, $\overline{T}yz =_\beta y$ and $\overline{F}yz =_\beta z$. Then δ is represented by $\lambda xyz.\, xyz$.

By specializing the "if then else" operator we can then easily represent all the connectives, as follows:

- *negation*
  Since

  $$\neg x = \begin{cases} F & \text{if } x = T \\ T & \text{if } x = F, \end{cases}$$

  negation is represented by $\lambda x.\, x\,\overline{F}\,\overline{T}$.

- *conjunction*
  Since
  $$x \wedge y = \begin{cases} y & \text{if } x = T \\ F & \text{if } x = F, \end{cases}$$
  conjunction is represented by $\lambda xy.\, xy\overline{F}$.
- *disjunction*
  Since
  $$x \vee y = \begin{cases} T & \text{if } x = T \\ y & \text{if } x = F, \end{cases}$$
  disjunction is represented by $\lambda xy.\, x\overline{T}y$.
- *implication*
  Since
  $$x \to y = \begin{cases} y & \text{if } x = T \\ T & \text{if } x = F, \end{cases}$$
  implication is represented by $\lambda xy.\, xy\overline{T}$. □

Corollary 15.1.13 *All truth-valued functions are representable in the Untyped Lambda Calculus*

Proof. By **??**, any truth-valued function is reducible to a combination of classical connectives. □

Having obtained representations for all the truth-valued functions, the next step is to investigate how faithful these representations are. Some of the usual properties of the connectives continue to hold for their representations: for example, that if $u =_\beta \overline{T}$ or $u =_\beta \overline{F}$, then $u \wedge \neg u =_\beta \overline{F}$ and $u \vee \neg u =_\beta \overline{T}$.

The next result shows, however, that it is impossible to represent implication in a way faithful to all of its classical properties.

Proposition 15.1.14 Curry's Paradox (Curry [1942]) *It is impossible to represent implication in the Untyped Lambda Calculus in such a way that it satisfies both Modus Ponens, i.e.*

$$\frac{u =_\beta \overline{T} \qquad (u \to v) =_\beta \overline{T}}{v =_\beta \overline{T}}$$

and any of the two following properties:

1. *if $u =_\beta (u \to v)$, then $u =_\beta \overline{T}$*
2. *$[u \to (u \to v)] \to (u \to v) =_\beta \overline{T}$.*

Proof. Given any term v, by the Fixed Point Theorem 12.2.2 there is a term u such that $u =_\beta (u \to v)$. If 1 holds, then both terms u and $u \to v$ are equal to $\overline{T}$. And if Modus Ponens holds, then v is equal to $\overline{T}$. But since v is *any* term, this means that all terms are β-equal, which they are not.

Similarly, given any term v, by the Fixed Point Theorem there is a term u such that $u =_\beta u \to (u \to v)$. Then

$$[u \to (u \to v)] \to (u \to v) \ =_\beta \ u \to (u \to v) \ =_\beta \ u.$$

If 2 holds, then all these three terms are equal to $\overline{T}$. And if Modus Ponens holds, then

$$\dfrac{u =_\beta \overline{T} \qquad \dfrac{u =_\beta \overline{T} \quad u \to (u \to v) =_\beta \overline{T}}{u \to v =_\beta \overline{T}}}{v =_\beta \overline{T}}\ ,$$

and again all terms are β-equal. □

Notice that both 1 and 2 are actually true of classical implication. To prove 1, suppose $\vdash_{\mathcal{N}} \alpha \leftrightarrow (\alpha \to \beta)$. If we assume α, we have $\alpha \to \beta$ by hypothesis, and hence β by Modus Ponens or $\to$-Elimination. This proves $\alpha \vdash_{\mathcal{N}} \beta$, and hence $\vdash_{\mathcal{N}} \alpha \to \beta$ by the Deduction Theorem or $\to$-Introduction, and $\vdash_{\mathcal{N}} \alpha$ by hypothesis again.

This informal reasoning is formalized by 2, which can easily be proved in either the Natural Deduction or the Sequent Systems:

$$\dfrac{\dfrac{\dfrac{[\alpha]^{(1)} \qquad \dfrac{[\alpha]^{(1)} \quad [\alpha \to (\alpha \to \beta)]^{(2)}}{\alpha \to \beta}}{\beta}}{\alpha^{(1)} \to \beta}}{[\alpha \to (\alpha \to \beta)]^{(2)} \to (\alpha \to \beta)} \qquad \dfrac{\dfrac{\dfrac{\alpha \vdash_{\mathcal{S}} \beta, \alpha \qquad \dfrac{\alpha \vdash_{\mathcal{S}} \beta, \alpha \quad \beta, \alpha \vdash_{\mathcal{S}} \beta}{\alpha \to \beta, \alpha \vdash_{\mathcal{S}} \beta}}{\alpha \to (\alpha \to \beta), \alpha \vdash_{\mathcal{S}} \beta}}{\alpha \to (\alpha \to \beta) \vdash_{\mathcal{S}} \alpha \to \beta}}{\vdash_{\mathcal{S}} [\alpha \to (\alpha \to \beta)] \to (\alpha \to \beta).}$$

Notice also that the proof of 15.1.14 does not use any particular representation of the truth-values, and it is thus a very general negative result regarding the possibility of a faithful representation of Classical Propositional Logic in the Untyped Lambda Calculus.

Exercise 15.1.15 Russell's Paradox à la Curry. *In any logical system including* $\to$*-Introduction and* $\to$*-Elimination, as well as set-theoretical Full Comprehension, every formula is provable.* (Hint: given any formula β, by Full Comprehension there is a set $R_\beta = \{x : x \in x \to \beta\}$. If we assume $R_\beta \in R_\beta$, we also have $R_\beta \in R_\beta \to \beta$ by definition of R_β, and hence β by $\to$-Elimination. This proves $R_\beta \in R_\beta \to \beta$ by $\to$-Introduction, and hence $R_\beta \in R_\beta$ by definition of R_β. Then β follows by $\to$-Elimination again.)

Thus negation is not necessary for Russell's Paradox, and implication is sufficient.

15.2 Undecidability

The positive fact of the representability of all recursive functions brings with it the negative by-product of the undecidability of all nontrivial problems on the Untyped Lambda Calculus. We first concentrate on some particularly interesting examples, and then prove a general result. The proof technique always consists of first translating the given problem into an appropriate function, and then proving that the translation is not recursive, by a standard application of the diagonal method (see Odifreddi [1989], Section II.2 for discussion of it).

Since the problems on the Lambda Calculus refer to terms, while the recursive functions are defined on numbers, to make the translation possible it is convenient to code terms by numbers in some canonical effective way. We do not specify any such coding, since the results we prove are independ of the details. We just refer to a generic effective enumeration $\{t_x\}_{x\in\omega}$ of all untyped λ-terms.

Examples

The proofs of the next results all mirror the standard proof of the existence of a recursively enumerable nonrecursive set.

Theorem 15.2.1 (Church [1936]) *β-equality is undecidable, in the sense that the set of pairs of β-equal λ-terms is not recursive.*

Proof. Consider the following function:

$$f(x) = \begin{cases} 1 & \text{if } t_x\overline{x} =_\beta \overline{0} \\ 0 & \text{otherwise.} \end{cases}$$

If β-equality were decidable, f would be recursive. By the Representation Theorem 15.1.7, there would then be a term F such that

$$F\overline{x} =_\beta \begin{cases} \overline{1} & \text{if } t_x\overline{x} =_\beta \overline{0} \\ \overline{0} & \text{otherwise.} \end{cases}$$

Then

$$F\overline{x} =_\beta \overline{0} \iff t_x\overline{x} \neq_\beta \overline{0}.$$

If $F = t_n$, i.e. if n is the code number of F, then

$$t_n\overline{x} =_\beta \overline{0} \iff t_x\overline{x} \neq_\beta \overline{0},$$

and hence

$$t_n\overline{n} =_\beta \overline{0} \iff t_n\overline{n} \neq_\beta \overline{0},$$

which is a contradiction. □

Since β-equality is finitely axiomatizable by 12.1.4, 12.1.5 and 12.1.6, by translating the conjunction of its axioms into the language of Classical Predicate Logic we obtain a formula with two free variables which represents β-equality, in the sense of being provable for two terms if and only if they are β-equal. From the undecidability of β-equality it thus follows that Classical Predicate Logic is undecidable, a result showing the *unsolvability of the Entscheidungsproblem* proposed by Hilbert. The proof just sketched is the original one given by Church [1936a].

An independent and simultaneous proof of the same result was given by Turing [1936] in terms of Turing machines, which were introduced by him for this purpose and later became the standard model of computability. The main result on which Turing based his proof was the *unsolvability of the Halting Problem*, which was a version of the next result.

Theorem 15.2.2 (Church [1936]) *Normalization is undecidable, in the sense that the set of λ-terms having a normal form is not recursive.*

Proof. Consider the following function:

$$f(x) = \begin{cases} 1 & \text{if } t_x\overline{x} \text{ has a normal form} \\ 0 & \text{otherwise.} \end{cases}$$

If normalization were decidable, then f would be recursive. By the Representation Theorem 15.1.7, there would be a term F such that

$$F\overline{x} =_\beta \begin{cases} \overline{1} & \text{if } t_x\overline{x} \text{ has a normal form} \\ \overline{0} & \text{otherwise.} \end{cases}$$

By definability of definition by cases, based on test on zero, there would be a term G such that

$$G\overline{x} =_\beta \begin{cases} \Delta\Delta & \text{if } t_x\overline{x} \text{ has a normal form} \\ \overline{0} & \text{otherwise.} \end{cases}$$

Then

$$G\overline{x} \text{ has a normal form } \Longleftrightarrow t_x\overline{x} \text{ does not have a normal form.}$$

If $G = t_n$, i.e. if n is the code number of G, then

$$t_n\overline{x} \text{ has a normal form } \Longleftrightarrow t_x\overline{x} \text{ does not have a normal form,}$$

and hence

$$t_n\overline{n} \text{ has a normal form } \Longleftrightarrow t_n\overline{n} \text{ does not have a normal form,}$$

which is a contradiction. □

Corollary 15.2.3 *The complexity of normalization is not recursive, in the sense that there is no recursive function bounding the number of steps needed to get a normal form of any λ-term having one, as a function of the length of the given term.*

Proof. Otherwise, to know whether a term has a normal form it would be enough to systematically produce all possible reductions of the given term up to the recursive bound, and see if a term in normal form is produced in the process. $\Box$

A different proof of the corollary can be obtained as in 11.4.1, based on the fact that every recursive function is representable.

The Scott-Curry Theorem

Rice's Theorem (see Odifreddi [1989], II.2.9) summarizes the undecidability results of Recursion Theory by showing that the only decidable sets of programs closed under functional equality (i.e. containing either none or all programs of the same function) are the trivial ones, namely the empty set and the set of all programs. The next result provides an analogue for the Untyped Lambda Calculus, with the role of programs played by the terms.

Theorem 15.2.4 (Scott [1963], Curry [1969]) *The only decidable sets of terms closed under β-equality (i.e. containing either none or all terms β-equal to any given term) are the trivial ones, namely the empty set and the set of all terms.*

Proof. Given a set of terms $\mathcal{A}$, consider the following function:

$$f(x) = \begin{cases} 1 & \text{if } t_x\overline{x} \in \mathcal{A} \\ 0 & \text{otherwise.} \end{cases}$$

If $\mathcal{A}$ were decidable, then f would be recursive. By the Representation Theorem 15.1.7, there would be a term F such that

$$F\overline{x} =_\beta \begin{cases} \overline{1} & \text{if } t_x\overline{x} \in \mathcal{A} \\ \overline{0} & \text{otherwise.} \end{cases}$$

If $\mathcal{A}$ is neither empty nor the set of all terms, there exist $u \in \mathcal{A}$ and $v \notin \mathcal{A}$. By definability of definition by cases, based on test on zero, there would be a term G such that

$$G\overline{x} =_\beta \begin{cases} v & \text{if } t_x\overline{x} \in \mathcal{A} \\ u & \text{otherwise.} \end{cases}$$

If $\mathcal{A}$ is closed under β-equality,

$$G\overline{x} \in \mathcal{A} \iff t_x\overline{x} \notin \mathcal{A}.$$

If $G = t_n$, i.e. if n is the code number of G, then

$$t_n\overline{x} \in \mathcal{A} \iff t_x\overline{x} \notin \mathcal{A},$$

and hence

$$t_n\overline{n} \in \mathcal{A} \iff t_n\overline{n} \notin \mathcal{A},$$

which is a contradiction. □

The Scott-Curry Theorem has an obvious generalization from sets of terms to n-ary relations, saying that the only decidable n-ary relations of terms closed under β-equality are the trivial ones, namely the empty set and the set of all n-tuples of terms.

It follows that *any nontrivial property of terms invariant under β-equality is undecidable*. Undecidability proofs are thus reduced to nontriviality proofs, which are usually quite immediate. For example, the previous examples of undecidability follow from the facts that some but not all pairs of terms are β-equal (15.2.1), and some but not all terms have a normal form (15.2.2). Similarly, being β-equal to a given fixed term is undecidable, because some but not all terms are β-equal to it.

The next corollary gives a different application.

Corollary 15.2.5 (Grzegorczyk) *The theory of β-equality is essentially undecidable, in the sense that it has no consistent decidable extension.*

Proof. Consider a consistent extension $\mathcal{T}$ of the theory of β-equality. Then the set of pairs of terms identified by $\mathcal{T}$ is a nontrivial set closed under β-equality, and hence is not decidable. □

æ

Chapter 16

Type Assignments

16.1 Simple Types

Definition 16.1.1 *A term t is* **simply typable** *if there is a context Γ and a type α such that $\Gamma \vdash_{\rightarrow} \alpha$.*

16.2 Intersection Types

Definition 16.2.1 *A term t is* **intersection typable** *if there is a context Γ and a type α such that $\Gamma \vdash_{\rightarrow\cap} \alpha$.*

Intersection types

Using simple types we can type exactly those untyped λ-terms that are well formed as typed λ-terms. Using intersection types, instead, we can type more untyped λ-terms, although not all.

A typical example of a λ-term which is intersection typable, but not simply typable, is xx. Indeed, it is enough to give x both types $\alpha \rightarrow \beta$ and α, to be able to give type β to xx:

$$\{x : (\alpha \rightarrow \beta) \cap \alpha\} \vdash xx : \beta.$$

A typical example of a λ-term which instead is not intersection typable is $(\lambda x.\, xx)(\lambda x.\, xx)$. Otherwise, there would be two equal types $\alpha \rightarrow \beta$ and α, which is impossible for intersection types as it was for simple types.

A characterization of strongly normalizable terms

The interest of intersection types comes from the fact that they provide a characterization of the strongly normalizable λ-terms of the Untyped Lambda Calculus.

Proposition 16.2.2 Strong Normalization (Pottinger [1980]) *Every intersection typable term is strongly normalizable.*

Proof. The proof of the Strong Normalization Theorem **??** goes through. □

As a warm-up we prove the following weaker result.

Proposition 16.2.3 (Coppo and Dezani [1980]) *Every term in normal form is intersection typable.*

Proof. We prove that if t is a term in normal form, it can be given *some* type. And if t does not begin by a λ, it can actually be given *any* type (this is needed in the induction step).

We proceed by induction on the following inductive definition of the terms in normal form:

1. variables are in normal form
2. if u is in normal form, then so is $\lambda x.\, u$
3. if t does not begin by a λ, and t and u are in normal form, then so is tu.

Obviously, variables can be given any type.

If u can be given a type β, there are two cases. If x occurs in u, then $\lambda x.\, u$ has type $\alpha_1 \cap \cdots \cap \alpha_n \to \beta$, where $\alpha_1, \ldots, \alpha_n$ are the types already assigned to x. If instead x does not occur in u, then u can be given any type α, and $\lambda x.\, u$ can be given type $\alpha \to \beta$, for any α.

If u is in normal form, by induction hypothesis it can be given some type α. If t is in normal form and does not begin by a λ, by induction hypothesis it can be given any type, in particular $\alpha \to \beta$. Then tu can be given type β, for any β. □

We now turn to the proof of the main result.

Theorem 16.2.4 Intersection Typability (Pottinger [1980]) *Every strongly normalizable term is intersection typable.*

Proof. We proceed by induction on the height of the reduction tree of a strongly normalizable term t.

As it can easily be checked by induction on the definition of terms, t (as any other term) can only be of one of two kinds, which we consider separately:

1. $t = \lambda x_1 \cdots \lambda x_n .\, x t_1 \cdots t_m$

   Since t is strongly normalizable, so are $t_1, \ldots, t_m$, with reduction trees of smaller height than t. Then, by induction hypothesis, they can all be typed, with typing assignments $\Gamma_i \vdash t_i : T_i$. By collecting all the assignments together, and giving x its old types plus the new one

   $$T_1 \to (T_2 \to \cdots (T_m \to X),$$

   for some X, we have for some Γ:

   $$\Gamma \vdash x t_1 \cdots t_m : X.$$

   By assigning types to $x_1, \ldots, x_n$, if they do not have them already, we can thus type t as well.

2. $t = \lambda x_1 \cdots \lambda x_n .\, (\lambda x.\, u) v t_1 \cdots t_m$

   This term reduces in one step to $\lambda x_1 \cdots \lambda x_n .\, u[x =: v] t_1 \cdots t_m$. Since t is strongly normalizable, so are $u[x =: v] t_1 \cdots t_m$ and v, with reduction trees of smaller height. Thus they can be typed from some assignment Γ, by the induction hypothesis. It remains to prove a Lifting Lemma showing that from Γ we can type $(\lambda x.\, u) v t_1 \cdots t_m$, and hence t. □

As announced, to finish the proof we need to prove the Lifting Lemma, and we start by the following special case.

Proposition 16.2.5 Lifting Lemma (special case) *If* $\Gamma \vdash v : \alpha$ *and* $\Gamma \vdash u[x := v] : \beta$, *then there is* γ *such that* $\Gamma \vdash v : \gamma$ *and* $\Gamma \cup \{x : \gamma\} \vdash u : \beta$. *Hence* $\Gamma \vdash (\lambda x.\, u) : \gamma \to \beta$ *and* $\Gamma \vdash (\lambda x.\, u) v : \beta$.

Proof. We proceed by induction on u.

1. *u is a variable*

   If $u = x$, then $\Gamma \cup \{x : \beta\} \vdash u : \beta$. Moreover $u[x := v] = v$, and $\Gamma \vdash v : \beta$ by hypothesis. Thus we can let $\gamma = \beta$.

   If $u = y \neq x$, then $u[x := v] = u$, and $\Gamma \vdash u : \beta$ by hypothesis. A fortiori, $\Gamma \cup \{x : \alpha\} \vdash u : \beta$. Since $\Gamma \vdash v : \alpha$ by hypothesis, we can let $\gamma = \alpha$.

2. $u = \lambda y.\, t$

   Since $u[x := v] = \lambda y.\, t[x := v]$, by hypothesis $\Gamma \vdash \lambda y.\, t[x := v] : \beta$. We will prove a Decomposition Lemma, according to which $\beta = \beta_1 \to \beta_2$ and $\Gamma \cup \{y : \beta_1\} \vdash t[x := v] : \beta_2$, for some β_1 and β_2. Then, by induction hypothesis, there is γ such that $\Gamma \cup \{y : \beta_1\} \vdash v : \gamma$ and

   $$\Gamma \cup \{y : \beta_1\} \cup \{x : \gamma\} \vdash t : \beta_2,$$

so that

$$\Gamma \cup \{y : \beta_1\} \cup \{x : \gamma\} \vdash \lambda y.\, t : \beta_1 \to \beta_2,$$

i.e.

$$\Gamma \cup \{y : \beta_1\} \cup \{x : \gamma\} \vdash u : \beta.$$

3. $u = t_1 t_2$

   Since $u[x := v] = (t_1[x := v])(t_2[x := v])$, by hypothesis $\Gamma \vdash (t_1[x := v])(t_2[x := v]) : \beta$. We will prove a Decomposition Lemma, according to which $\Gamma \vdash t_1[x := v] : \delta \to \beta$ and $\Gamma \vdash t_2[x := v] : \delta$, for some δ. Then, by induction hypothesis, there is γ such that $\Gamma \vdash v : \gamma$ and

   $$\Gamma \cup \{x : \gamma\} \vdash t_1 : \delta \to \beta \quad \text{and} \quad \Gamma \cup \{x : \gamma\} \vdash t_2 : \delta,$$

   so that

   $$\Gamma \cup \{x : \gamma\} \vdash t_1 t_2 : \beta,$$

   i.e.

   $$\Gamma \cup \{x : \gamma\} \vdash t : \beta. \quad \Box$$

To conclude the proof of the Lifting Lemma, and hence of the Intersection Typability Theorem, we still need to prove the following result.

Proposition 16.2.6 Decomposition Lemma. *If $\Gamma \vdash t : \beta$, and β is not an intersection type, then:*

1. *if t is a variable, then Γ contains $x : \beta$*
2. *if $t = \lambda x.\, t_1$, then $\beta = \beta_1 \to \beta_2$ and $\Gamma \cup \{x : \beta_1\} \vdash t_1 : \beta_2$*
3. *if $t = t_1 t_2$, then $\Gamma \vdash t_2 : \delta$ and $\Gamma \vdash t_1 : \delta \to \cdots \cap \beta \cap \cdots$*

Proof. The result is immediate by induction on the type assignment rules, since the corresponding system in cut-free. $\Box$

We can now turn to the general case of the Lifting Lemma.

Proposition 16.2.7 Lifting Lemma (general case) *If $\Gamma \vdash v : \alpha$ and $\Gamma \vdash u[x := v]t_1 \cdots t_m : \beta$, then $\Gamma \vdash (\lambda x.\, u)vt_1 \cdots t_m : \beta$.*

Proof. By induction on n and β.

If $\beta = \beta = \beta_1 \cap \beta_2$ is an intersection type, and $\Gamma \vdash u[x := v] : \beta_1 \cap \beta_2$, then $\Gamma \vdash u[x := v] : \beta_1$ and $\Gamma \vdash u[x := v] : \beta_2$. By induction hypothesis on β, $\Gamma \cap \{x : \gamma\} \vdash u : \beta_1$ and $\Gamma \cap \{x : \gamma\} \vdash u : \beta_2$, so that $\Gamma \cap \{x : \gamma\} \vdash u : \beta_1 \cap \beta_2$.

If β is not an intersection type, then we proceed by induction on n. We already proved the special case $n = 0$ above. Suppose now $\Gamma \vdash u[x := v]t_1 \cdots t_n t_{n+1} : \beta$.

By the Decomposition Lemma, there is δ such that $\Gamma \vdash u[x := v]t_1 \cdots t_n : \delta \rightarrow \cdots \cap \beta \cap \cdots$ and $\Gamma \vdash t_{n+1} : \delta$. Then

$$\Gamma \vdash (\lambda x.\, u)vt_1 \cdots t_n : \delta \rightarrow \cdots \cap \beta \cap \cdots$$

by the induction hypothesis, hence

$$\Gamma \vdash (\lambda x.\, u)vt_1 \cdots t_n t_{n+1} : \cdots \cap \beta \cap \cdots$$

and in particular

$$\Gamma \vdash (\lambda x.\, u)vt_1 \cdots t_n t_{n+1} : \beta. \quad \Box$$

16.3 ω-Intersection Types

In the simple typing system the same term may be given different types, by different assignments of types to its variables. For example, the identity $\lambda x.\, x$ can be given infinitely many types, namely $\alpha \rightarrow \alpha$ for any type α, because $\{x : \alpha\} \vdash \lambda x.\, x : \alpha \rightarrow \alpha$.

If we identify terms which are α-convertible, then the same term may actually have different types, since e.g. from the assignment $x : \alpha \rightarrow \alpha$ and $y : \alpha$ we get $(\lambda x.\, x) : (\alpha \rightarrow \alpha) \rightarrow (\alpha \rightarrow \alpha)$ and $(\lambda y.\, y) : \alpha \rightarrow \alpha$. In particular, the term $\mathbf{I} = \lambda x.\, x =_\alpha \lambda y.\, y$ has the two types $\alpha \rightarrow \alpha$ and α, and hence $\mathbf{II} : \alpha \rightarrow \alpha$. Thus $\mathbf{II}$ is typable. But xx is not, because the same variable can only be assigned one type at a time.

In particular, the simple typing system is not closed under β-reductions, i.e. it is not true in general that if both $u[x := v]$ and v are typable, then so is $(\lambda x.\, u)v$. For example, $(\lambda x.xx)\mathbf{I}$ is not, although $\mathbf{I}$ and $\mathbf{II}$ are. What fails is that $u[x :=]$ may contain occurrences of the same term v (modulo α-equivalence) with different types. But with intersection types we can assign x *all* those types, and then $(\lambda x.u)v$ becomes typable. The remaining case, i.e. that x does not occur free in u, is trivial, since it is enough to assign x the type of v, which exists by hypothesis.

But even with conjunction types we do not have that if $t \rightarrow_\beta t_1$, and t_1 is typable, then so is t. For example, $(\lambda xy.\, y)(\Delta\Delta)\mathbf{I}$ β-reduces to $\mathbf{I}$, which is typable, but $\Delta\Delta$ is not typable. To avoid this, and get a typing system closed under β-reduction, we introduce a universal type ω.

Definition 16.3.1 (Sallé [1978], Coppo, Dezani and Venneri [1981]) *Every term t has type ω in any context, i.e.*

$$\frac{}{\Gamma \vdash t : \omega}$$

Definition 16.3.2 *A term t is* **ω-intersection typable** *if there is a context Γ and a type $\alpha \neq \omega$ such that $\Gamma \vdash_{\rightarrow\cap\omega} t : \alpha$.*

A characterization of normalizable terms

Theorem 16.3.3 ω-Intersection Typability (Coppo, Dezani and Venneri [1981]) *Every normalizable term is ω-intersection typable.*

Proof. Since terms in normal form are intersection typable by 16.2.3, it is enough to show that ω-interesection types are inherited upwards by β-reduction, i.e. that if $t \to_\beta t_1$ and t_1 has ω-intersection type α, then so does t.

Most of the work has already been done in the proof of the special case 16.2.5 of the Lifting Lemma, which says that if v is typable and $u[x := v]$ has type β, then so does $(\lambda x.\, u)v$. Notice that the hypothesis that v is typable is always satisfied with ω-intersection types, since v always has type ω. The rest follows by induction on t and α, as follows.

If $\alpha = \alpha_1 \cap \alpha_2$, the result is trivial. And if $\alpha \neq \omega$ is not an intersection type, then we proceed by induction on t:

- *t is a variable*

  Then no β-reduction is possible inside it.

- $t = \lambda x.\, u$

  Then $t_1 = \lambda x.\, u_1$, so $\alpha = \alpha_1 \to \alpha_2$. By the Decomposition Lemma, $\Gamma \cup \{x : \alpha_1\} \vdash u_1 : \alpha_2$, and by induction hypothesis $\Gamma \cup \{x : \alpha_1\} \vdash u : \alpha_2$. Then $\Gamma \vdash (\lambda x.\, u) : \alpha$.

- $t = uv$

  If $t_1 = u_1 v$ or $t_1 = uv_1$, then we can use again the Decomposition Lemma and the induction hypothesis. If instead u begins by a λ and t_1 is a reduct, we can use the special case of the Lifting Lemma already proved. □

Theorem 16.3.4 Normalization (Coppo, Dezani and Venneri [1981]) *Every ω-intersection typable term is normalizable.*

Proof. The proof is a variation of the proof of the Strong Normalization Theorem **??**. We define a class $|calC = \bigcup_\alpha \mathcal{C}_\alpha$ by induction on α, as follows:

- if α is atomic or $\alpha = \omega$, then

$$t \in \mathcal{C}_\alpha \;\Leftrightarrow\; t \text{ is normalizable (by leftmost reductions)}$$

- $t \in \mathcal{C}_{\alpha\to\beta}$ if and only if $(\forall u \in \mathcal{C}_\alpha)(tu \in \mathcal{C}_\beta)$
- $t \in \mathcal{C}_{\alpha\cap\beta}$ if and only if $t \in \mathcal{C}_\alpha$ and $t \in \mathcal{C}_\beta$, i.e. $\mathcal{C}_{\alpha\cap\beta} = \mathcal{C}_\alpha \cap \mathcal{C}_\beta$.

Then, in a way similar to (but simpler than) the proof of the Strong Normalization Theorem for the Typed Lambda Calculus, we have:

1. *$\mathcal{C}$ contains only terms which are normalizable (by leftmost reductions)*

   As usual, we prove by simultaneous induction that:

   - if $u_1, \ldots, u_n$ are normalizable (by leftmost reductions), then $xu_1 \cdots u_n \in \mathcal{C}$
   - if $t \in \mathcal{C}$, then t is normalizable (by leftmost reductions).

   The only non trivial case is the type $\alpha \to \beta$. But if $t \in \mathcal{C}_{\alpha\to\beta}$, it is enough to choose $x \in \mathcal{C}_\alpha$ not occurring in t. Then $tx \in \mathcal{C}$ and is normalizable (by leftmost reductions).

   If t does not begin by a λ, then reductions are possible only within t, and hence t is normalizable (by leftmost reductions).

   If $t = \lambda y.\, t_1$, then the normalization (by leftmost reductions) of tx produces a normalization of $t_1[y := x]$ (which is obtained by the leftmost reduction), and hence of t_1 itself (since y has only changed name) and of t.

2. *$\mathcal{C}$ contains every ω-intersection typable term*

   As usual, we prove that if $v_1, \ldots, v_n \in \mathcal{C}$, then $t[x_1 := v_1, \ldots, x_n := v_n] \in \mathcal{C}$.

   The only interesting case is $t = \lambda x.\, u$. But then

   $$(\lambda x.\, u)[x_1 := v_1, \ldots, x_n := v_n] = \lambda x.\, (u[x_1 := v_1, \ldots, x_n := v_n]),$$

   which is in $\mathcal{C}$ if and only if so is

   $$\lambda x.\, (u[x_1 := v_1, \ldots, x_n := v_n])v,$$

   for every $v \in \mathcal{C}$. But the previous term reduces to

   $$u[x_1 := v_1, \ldots, x_n := v_n, x := v].$$

   We thus need to prove that

   $$\text{if } a[x := v] \in \mathcal{C} \text{ and } v \in \mathcal{C}, \text{ then } (\lambda x.\, a)v \in \mathcal{C}.$$

   By going to atomic types or ω, by a sequence of appropriate applications, we have to show this term is normalizable (by leftmost reductions). But this is trivial, since $(\lambda x.\, a)v$ reduces (by leftmost reductions) to $a[x := v]$. □

16.4 Filter Models

Terms as collections of types

$[\![t]\!] = \{\alpha : t \text{ can be assigned type } \alpha\}$: model of Untyped Lambda Calculus
relationships with D_∞

Types as collections of terms

$[\![\alpha]\!] = \{t : t \text{ can be assigned type } \alpha\}$: model of intuitionistic logic?

Part E

Intuitionistic Propositional Calculus

Chapter 17

Intuitionistic Propositional Calculus

Until now we have only looked at the two connectives of implication and conjunction, in various ways. We now turn to a fuller presentation of the Intuitionistic Propositional Calculus, and consider other connectives.

Section 1 deals with *disjunction*, which is symmetric to conjunction from an algebraic point of view. The only significant failure of symmetry, i.e. the existence of a greatest, but not of a least element in the 'term model' of formulas modulo provable equivalence, is remedied in Section 2 by the introduction of a constant for *falsity*, which can then be used to define *negation*.

17.1 Disjunction

Implicational Calculus with Disjunction

We extend Implicational Calculus as follows:

1. the **language** has an added connective $\vee$ (*disjunction*)

2. the definition of **formulas** has an added inductive clause, i.e.

    - if α and β are formulas, so is $(\alpha \vee \beta)$.

To increase readability, some parentheses can be omitted according to the precedence rule: *disjunction over implication*. When also conjunction is present, there is no precedence rule between $\wedge$ and $\vee$.

The goal of this section is to determine which of the formulas of the Implicational Calculus with Disjunction can be considered 'true', when the connective $\vee$ is intuitively taken as representing 'disjunction'.

Following the blueprint of Chapters 1–5, we introduce different but equivalent analyses. We continue to use the same symbols $\vdash_{\mathcal{N}}$, $\vdash_{\mathcal{H}}$, $\vdash_{\mathcal{S}}$, $\vdash_{\mathcal{T}}$, $\models_i$ and $\models_a$, but they now refer to the extended system with implication, conjunction and disjunction.

Natural Deduction

To justify the rules for $\vee$ we go back to the original motivation of Natural Deduction, as a system whose rules allow us to continue a given proof from assumptions.

Since $\vee$ is going to be interpreted as a disjunction, to *prove* $\alpha \vee \beta$ we have to prove α or β. This suggests the rules

$$\frac{\alpha}{\alpha \vee \beta} \qquad \text{and} \qquad \frac{\beta}{\alpha \vee \beta.}$$

The way to *use* disjunctions is suggested by a venerable principle already known to the Greeks, the socalled *proof by cases*: since a proof of $\alpha \vee \beta$ actually codes a proof of either α or β, a proof of γ from $\alpha \vee \beta$ can be completed by any proof of either α or β. But when proving γ from $\alpha \vee \beta$ we do not yet know whether the proof will be completed by a proof of α or β, and we must be ready for both possibilities. A proof from $\alpha \vee \beta$ will then actually code two proofs, one ($\mathcal{D}^{\alpha}_{\gamma}$) from α and one ($\mathcal{D}^{\beta}_{\gamma}$) from β.

As in the case of implication, α and β may appear in packets of occurrences as hypotheses in $\mathcal{D}^{\alpha}_{\gamma}$ and $\mathcal{D}^{\beta}_{\gamma}$, respectively. But after merging the two proofs into a single one from $\alpha \vee \beta$ we do not need to consider them as hypotheses anymore, since the information that a proof of α or β would complete a proof of γ is already contained in the fact that $\alpha \vee \beta$ now appears as an hypothesis: the packets of occurrences of α and β may thus be discharged in their respective proofs.

With the usual notations for discharge, we can thus picture the $\vee$-elimination rule as the step

$$\text{from} \quad \begin{array}{c}\Gamma, \alpha\\ \mathcal{D}^{\alpha}_{\gamma}\\ \gamma\end{array} \quad \text{and} \quad \begin{array}{c}\Gamma, \beta\\ \mathcal{D}^{\beta}_{\gamma}\\ \gamma\end{array} \quad \text{to} \quad \dfrac{\alpha \vee \beta \qquad \begin{array}[b]{c}\Gamma, [\alpha]^{(1)}\\ \mathcal{D}^{\alpha}_{\gamma}\\ \gamma\end{array} \qquad \begin{array}[b]{c}\Gamma, [\beta]^{(1)}\\ \mathcal{D}^{\beta}_{\gamma}\\ \gamma\end{array}}{\gamma^{(1)}.}$$

Definition 17.1.1 (Gentzen [1935]) *The relation $\vdash_{\mathcal{N}}$ defined in 1.1.1 and 4.1.1 is extended to disjunction as follows:*

6. $\vee$**-Introduction**. *If any of α and β is deducible from Γ, then so is $\alpha \vee \beta$:*

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha}{\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta} \qquad \textit{and} \qquad \frac{\Gamma \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta.}$$

7. $\vee$**-Elimination**. *If γ is deducible from Γ and both α and β separately, then it is also deducible from Γ and $\alpha \vee \beta$:*

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta \qquad \Gamma, \alpha \vdash_{\mathcal{N}} \gamma \qquad \Gamma, \beta \vdash_{\mathcal{N}} \gamma}{\Gamma \vdash_{\mathcal{N}} \gamma.}$$

In an application of the $\vee$-elimination rule, $\alpha \vee \beta$ is called the *major premise* and γ the *minor premise*.

The $\vee$-elimination rule defines a proof of γ from $\alpha \vee \beta$ as a pair of incomplete proofs of γ, one from a packet of assumptions α and another one from a packet of assumptions β, waiting for a completion. The $\vee$-introduction rule allows the completion of such a proof, whenever we have a proof of α or β. Taken together, the two rules combine in producing a proof of γ, for example as follows:

$$\frac{\dfrac{\begin{array}{c} \Gamma \\ \mathcal{D}_\alpha \\ \alpha \end{array}}{\alpha \vee \beta} \qquad \begin{array}{c} \Gamma, [\alpha] \\ \mathcal{D}^\alpha_\gamma \\ \gamma \end{array} \qquad \begin{array}{c} \Gamma, [\beta] \\ \mathcal{D}^\beta_\gamma \\ \gamma \end{array}}{\gamma.}$$

The occurrence of $\alpha \vee \beta$ in such a proof is called a **maximum** (relative to $\vee$).

A more direct way of getting to a proof of γ is obviously to forget about $\mathcal{D}^\beta_\gamma$, and the step from the above to the following:

$$\begin{array}{c} \Gamma \\ \mathcal{D}_\alpha \\ \Gamma, \quad \alpha \\ \mathcal{D}^\alpha_\gamma \\ \gamma \end{array}$$

is called a **maximum elimination**. A symmetric maximum elimination can be obtained by working on β.

We would like to say, as in Chapters 1 and 4, that a proof is in normal form if it has no maxima relative to $\rightarrow$, $\wedge$ or $\vee$, but this notion is not strong enough. Indeed, one nice feature of normal proofs is the Subformula Property, and with the proposed notion of normal proof this would fail. Consider e.g. the following proof, which is normal in the sense just given:

$$\dfrac{\dfrac{\alpha \vee \alpha \qquad \dfrac{[\alpha] \quad [\alpha]}{\alpha \wedge \alpha} \qquad \dfrac{[\alpha] \quad [\alpha]}{\alpha \wedge \alpha}}{\alpha \wedge \alpha}}{\alpha.}$$

Here α is proved from the only assumption $\alpha \vee \alpha$, but the proof uses the formula $\alpha \wedge \alpha$, which is not a subformula of either the assumption $\alpha \vee \alpha$ or the conclusion α. We can say that $\alpha \wedge \alpha$ is a **delayed maximum**, in the sense that it is first introduced in the subproof

$$\frac{\alpha \qquad \alpha}{\alpha \wedge \alpha}$$

and then eliminated in the subproof

$$\frac{\alpha \wedge \alpha}{\alpha,}$$

with a procrastination due to the fact that the formula $\alpha \wedge \alpha$ is dragged along in a whole **segment of occurrences**, without any direct action on it.

The main point of the following development is that *the properties of $\mathcal{N}$ are preserved by the extension to $\vee$, provided we substitute the notion of 'occurrence' of a formula with that of 'segment of occurrences'.*

In particular, we can extend the notion of **normal proof** as follows: a proof is normal if there is no delayed maximum, i.e. no segment of occurrences of a given formula, starting with the consequence of an introduction rule, and ending with the (major) premise of an elimination rule (of the same connective).[1] Maxima in the old sense are special cases of the new ones, when the segments have length 1.

We can also extend the notion of descending path of a proof (p. 11) by looking at sequences of segments, as opposed to sequences of formulas. The intention is that in a normal proof, except for the repetitions of formulas in any given segment, a path should first go through eliminations and then through introductions. We can then define a **descending path** as any maximal sequence of consecutive formulas that:

1. does not go through minor premises of $\rightarrow$-eliminations

2. when (descending in $\mathcal{D}_{\alpha\vee\beta}$) it reaches a major premise $\alpha\vee\beta$ of a $\vee$-elimination

$$\frac{\begin{array}{c}\Gamma\\ \mathcal{D}_{\alpha\vee\beta}\\ \alpha\vee\beta\end{array} \qquad \begin{array}{c}\Gamma,[\alpha]\\ \mathcal{D}^{\alpha}_{\gamma}\\ \gamma\end{array} \qquad \begin{array}{c}\Gamma,[\beta]\\ \mathcal{D}^{\beta}_{\gamma}\\ \gamma\end{array}}{\gamma,}$$

[1] An even stronger notion of normal proof (for which the results proved below would still hold) requires also no *redundant application of a $\vee$-elimination rule*, i.e. one

$$\frac{\begin{array}{c}\Gamma\\ \mathcal{D}_{\alpha\vee\beta}\\ \alpha\vee\beta\end{array} \qquad \begin{array}{c}\Gamma,[\alpha]\\ \mathcal{D}^{\alpha}_{\gamma}\\ \gamma\end{array} \qquad \begin{array}{c}\Gamma,[\beta]\\ \mathcal{D}^{\beta}_{\gamma}\\ \gamma\end{array}}{\gamma}$$

in which either α is not discharged in $\mathcal{D}^{\alpha}_{\gamma}$, or β is not discharged in $\mathcal{D}^{\beta}_{\gamma}$.

then (if it continues) it jumps to a discharged occurrence of α in $\mathcal{D}^\alpha_\gamma$ or of β in $\mathcal{D}^\beta_\gamma$.

The reasons for the novel clause is that we only want to consider full segments, and γ might be part of a segment that starts above it (in $\mathcal{D}^\alpha_\gamma$ or $\mathcal{D}^\beta_\gamma$): if we went from $\alpha \vee \beta$ to γ we would lose a part of that segment. Similarly for the maximality condition.

With the new notions of normal proof and of descending path we can prove, as usual, the following result.

Proposition 17.1.2 Structure of Normal Proofs (Prawitz [1965]) *For a normal proof of $\mathcal{N}$ the following hold:*

1. **Elimination-Introduction Separation**. *Disregarding repetitions of formulas, any descending path consists of two (possibly empty) parts: a first (upper) one going only through elimination rules, and a second (lower) one going only through introduction rules.*

2. **Subformula Property**. *Any formula occurring in the proof is a subformula of either an undischarged assumption or the conclusion.*

Proof. Each descending path is a sequence $S_1, \ldots, S_n$ of segments, and each segment S_i consists of a sequence of occurrences of a same formula. By definition of path, the last occurrence of S_i is the premise of a rule R_i, and the first occurrence of S_{i+1} is either the consequence of the same rule R_i, or it comes from a jump as in clause 2 of the definition of path (in which case R_i must be a $\vee$-elimination).

Suppose R_i is an introduction rule (in particular, *not* a $\vee$-elimination), and R_{i+1} is an elimination rule. Then the first occurrence of S_{i+1} is the consequence of an introduction rule, while the last occurrence is a premise of an elimination rule. Moreover, the latter occurrence must be the *major* premise of such a rule, since the only possible cases are: $\rightarrow$-elimination (and a path does not go through a minor premise of $\rightarrow$-eliminations), $\wedge$-elimination (for which there are no minor premises), or $\vee$- elimination (and a minor premise of a $\vee$-elimination cannot be the last occurrence of a segment). Then the same connective must be introduced by the rule of which the first occurrence of S_{i+1} is the consequence, and eliminated by the rule of which the last occurrence of S_{i+1} is the major premise. This means that the formula occurring in S_{i+1} is a delayed maximum, which is impossible if the proof is normal. Thus R_i and R_{i+1} cannot be, respectively, an introduction and an elimination rule. This proves the first part.

The second part is proved as in 1.1.2, by noting that: if R_i and R_{i+1} are both elimination rules (or $i = 1$), then the formula occurring in S_{i+1} is a subformula of the one occurring in S_i; and if R_i and R_{i+1} are both introduction rules (or $i = n$), then the formula occurring in S_i is a subformula of the one occurring in S_{i+1}. □

Also the Weak Normalization Theorem continues to hold.

Theorem 17.1.3 Weak Normalization (Prawitz [1965]) *Every proof can be transformed into a normal proof, by means of an appropriate sequence of maxima eliminations.*

Proof. As in the case of maxima relative to →, the elimination of a maximum

$$\dfrac{\dfrac{\begin{matrix}\Gamma\\ \mathcal{D}_\alpha\\ \alpha\end{matrix}}{\alpha\vee\beta} \qquad \begin{matrix}\Gamma,[\alpha]\\ \mathcal{D}^\alpha_\gamma\\ \gamma\end{matrix} \qquad \begin{matrix}\Gamma,[\beta]\\ \mathcal{D}^\beta_\gamma\\ \gamma\end{matrix}}{\gamma.} \qquad \text{into} \qquad \begin{matrix} & \Gamma\\ & \mathcal{D}_\alpha\\ \Gamma, & \alpha\\ & \mathcal{D}^\alpha_\gamma\\ & \gamma\end{matrix}$$

in a proof $\mathcal{D}$ can have the following two bad effects:

- it can increase the total number of maxima, since it reproduces $\mathcal{D}_\alpha$ (and hence all maxima occurring in it) above every occurrence of α in the package of assumptions used in $\mathcal{D}^\alpha_\gamma$, and there may be many such occurrences
- it can introduce a new maximum, if $\mathcal{D}_\alpha$ ends with an introduction of a connective, by turning into a maximum every occurrence of α below which $\mathcal{D}^\alpha_\gamma$ continues with an elimination of the same connective.[2]

As in 1.1.3, we only need to extend the notion of complexity to take care of the case of disjunction as well, by adding to the definition of **degree** of a formula the following clause:

- the degree of $\alpha \wedge \beta$ is 1 plus the greatest of the degrees of α and β.

The obvious attack is now to mimic the proof of 1.1.3: we eliminate, at every step, *a maximum of greatest degree, and such that no occurrence of a maximum of greatest degree occurs above it*. This works fine, as usual, if the maximum that is eliminated is not delayed. But we now also have to take care of (segments of) delayed maxima: in this case, as we could guess from the interlocutory steps in the proof of the Cut Elimination Theorem for $\mathcal{S}$, we have to reduce the length of the relative segments, by modifying the original proof.

Consider a segment of length > 1 corresponding to a maximum γ of greatest degree: the next to last occurrence of γ in it must be a minor premise of a $\vee$-elimination rule (since this is the only rule that drags formulas along, and can thus produce delayed maxima), and the last occurrence of γ must be the major premise

[2] Notice that, unlike in the case of maxima w.r.t. →, nothing bad can occur relatively to γ: if $\mathcal{D}^\alpha_\gamma$ ends with an introduction of a connective and $\mathcal{D}$ continues below γ with an elimination of the same connective, then γ is a delayed maximum already in $\mathcal{D}$, and actually the length of the relative segment decreases by one.

of an elimination rule (since γ is a delayed maximum, and thus it is introduced at the step preceeding its first occurrence in the segment, and eliminated at the step following its last occurrence). Thus the proof $\mathcal{D}$ continues after the last occurrence of γ with an application of one of the elimination rules. We only consider the case of $\to$-elimination, since the other cases are similar.

If γ is the major premise of a $\to$-introduction, then $\gamma = \gamma_1 \to \gamma_2$ and the proof looks like:

$$\dfrac{\begin{matrix}\mathcal{D}_4\\ \gamma_1\end{matrix} \qquad \dfrac{\begin{matrix}\mathcal{D}_1\\ \alpha\vee\beta\end{matrix} \qquad \begin{matrix}\mathcal{D}_2\\ \gamma_1\to\gamma_2\end{matrix} \qquad \begin{matrix}\mathcal{D}_3\\ \gamma_1\to\gamma_2\end{matrix}}{\gamma_1\to\gamma_2}}{\begin{matrix}\gamma_2\\ \mathcal{D}_5.\end{matrix}}$$

We transform it into the following:

$$\dfrac{\begin{matrix}\mathcal{D}_1\\ \alpha\vee\beta\end{matrix} \qquad \dfrac{\begin{matrix}\mathcal{D}_4\\ \gamma_1\end{matrix} \qquad \begin{matrix}\mathcal{D}_2\\ \gamma_1\to\gamma_2\end{matrix}}{\gamma_2} \qquad \dfrac{\begin{matrix}\mathcal{D}_4\\ \gamma_1\end{matrix} \qquad \begin{matrix}\mathcal{D}_3\\ \gamma_1\to\gamma_2\end{matrix}}{\gamma_2}}{\begin{matrix}\gamma_2\\ \mathcal{D}_5.\end{matrix}}$$

This transformation produces a new copy of $\mathcal{D}_4$ (and hence of all maxima occurring in it), and it could thus increase the total number of occurrences of maxima of greatest degree. To avoid this, as well as a similar problem in the case of $\vee$-elimination, we reformulate the condition for the normalization procedure as follows.

At any step, choose *an occurrence of a maximum which:*

1. *has greatest degree*

2. *if it belongs to a segment of length 1, then no maximum of greatest degree occurs above it*

3. *if it belongs to a segment of length greater than 1, then it is the last occurrence of the segment, and no maximum of greatest degree occurs above the minor premise(s) of the elimination rule of which the given occurrence is the major one.*

Then the elimination of the chosen occurrence decreases by one the total number of occurrences of maxima of greatest degree. □

Formally, the proof of the Weak Normalization Theorem is still by ω^2*-induction*, on the pair

(greatest degree, number of occurrences of maxima with greatest degree.)

Similarly, we could do induction not on the number of occurrences of maxima of greatest degree, but on the sum of the lengths of segments of maxima with greatest degree, since each step of the proof either decreases the length of one of these segments with length greater than 1, or it eliminates one of the segments with length 1.

The Normalization Theorem and the last part of the Subformula Property show that no new formulas of the Implicational Calculus with Conjunction can be proved in the extended system with disjunction: if $\vee$ does not occur in the premises or in the conclusion of a normal proof, then it does not occur at all in the proof. In technical terms, *the system with implication, conjunction and disjunction is a conservative extension of the system with implication and conjunction alone*.

Hilbert systems

As in the case of conjunction, we add axioms relative to the new connective $\vee$ mimicking $\vee$-introduction and $\vee$-elimination, and keep Modus Ponens as the only rule. The definition of $\vdash_{\mathcal{H}}$ is thus unchanged, and the Deduction Theorem is still valid as usual.

Theorem 17.1.4 Equivalence of Hilbert Systems and Natural Deduction (Gentzen [1935]) *If $\mathcal{H}$ is any Hilbert system whose theorems include 1–6 of 1.2.3 and 4.1.4 and, for any α, β and γ, the following:*

7. $\alpha \to \alpha \vee \beta$

8. $\beta \to \alpha \vee \beta$

9. $(\alpha \to \gamma) \to [(\beta \to \gamma) \to (\alpha \vee \beta \to \gamma)]$,[3]

then for any Γ and β:

$$\Gamma \vdash_{\mathcal{H}} \beta \ \Leftrightarrow\ \Gamma \vdash_{\mathcal{N}} \beta.$$

Proof. The right-to-left direction is obtained by induction on the definition of $\vdash_{\mathcal{N}}$. The only cases not dealt with in 1.2.2 and 4.1.4 are the ones relative to $\vee$, which we now treat.

Suppose $\Gamma \vdash_{\mathcal{H}} \alpha$. Then we get $\Gamma \vdash_{\mathcal{H}} \alpha \vee \beta$ by inserting the given proof of α from Γ above its occurrence in the following partial proof (from assumption 7):

$$\frac{\alpha \qquad \alpha \to \alpha \vee \beta}{\alpha \vee \beta.}$$

[3] Equivalently, and a bit more legibly, we could take

$$[(\alpha \to \gamma) \wedge (\beta \to \gamma)] \to (\alpha \vee \beta \to \gamma).$$

The advantage of the form above is that it does not use $\wedge$, and thus it can be used for an axiomatization of the fragment of propositional logic with → and ∨ alone.

Similarly from $\Gamma \vdash_{\mathcal{H}} \beta$, using 8.

Suppose

$$\Gamma \vdash_{\mathcal{H}} \alpha \vee \beta \qquad \Gamma, \alpha \vdash_{\mathcal{H}} \gamma \qquad \Gamma, \beta \vdash_{\mathcal{H}} \gamma.$$

By the Deduction Theorem we have $\Gamma \vdash_{\mathcal{H}} \alpha \to \gamma$ and $\Gamma \vdash_{\mathcal{H}} \beta \to \gamma$. Then we get $\Gamma \vdash_{\mathcal{H}} \gamma$ by inserting the given proofs of $\alpha \vee \beta$, $\alpha \to \gamma$ and $\beta \to \gamma$ from Γ above their occurrences in the following partial proof (from assumption 9):

$$\dfrac{\alpha\vee\beta \qquad \dfrac{\beta\to\gamma \qquad \dfrac{\alpha\to\gamma \qquad (\alpha\to\gamma)\to[(\beta\to\gamma)\to(\alpha\vee\beta\to\gamma)]}{(\beta\to\gamma)\to(\alpha\vee\beta\to\gamma)}}{\alpha\vee\beta\to\gamma}}{\gamma.}$$

For the left-to-right direction, we only need to show that 7, 8 and 9 are provable in Natural Deduction. 9 is proved by

$$\dfrac{\dfrac{\dfrac{\dfrac{[\alpha\vee\beta]^{(2)} \qquad \dfrac{[\alpha]^{(1)} \quad [\alpha\to\gamma]^{(4)}}{\gamma} \qquad \dfrac{[\beta]^{(1)} \quad [\beta\to\gamma]^{(3)}}{\gamma}}{\gamma^{(1)}}}{(\alpha\vee\beta)^{(2)}\to\gamma}}{(\beta\to\gamma)^{(3)}\to[\alpha\vee\beta\to\gamma]}}{(\alpha\to\gamma)^{(4)}\to[(\beta\to\gamma)\to(\alpha\vee\beta\to\gamma)].}$$

7 is proved by

$$\dfrac{\dfrac{[\alpha]^{(1)}}{\alpha\vee\beta}}{\alpha^{(1)}\to\alpha\vee\beta.}$$

8 is proved similarly. □

Sequents

As already for conjunction, the extension of the Sequent System to disjunction is unproblematic.

Definition 17.1.5 (Gentzen [1935]) *The relation $\vdash_{\mathcal{S}}$ defined in 1.3.1 and 4.1.5 is extended to disjunction as follows:*

6. **$\vee$-Introduction on the right**. *If α or β are deducible from Γ, then so is $\alpha \vee \beta$:*

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha}{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta} \quad \textit{and} \quad \frac{\Gamma \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta.}$$

7. $\vee$**-Introduction on the left**. *If γ is deducible from Γ and both α and β separately, then it is deducible from Γ and $\alpha \vee \beta$:*

$$\frac{\Gamma, \alpha \vdash_{\mathcal{S}} \gamma \quad \Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \vee \beta \vdash_{\mathcal{S}} \gamma.}$$

Unlike the case of conjunction, we cannot rephrase the two rules of $\vee$-introduction on the right as a single rule

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha, \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta,}$$

because the antecedent would have *two* consequences as opposed to one. In Chapter **??** we will see how Classical Logic can be defined by a sequent system that differs from the one considered here precisely in this aspect, of allowing more than a single consequence. But this will require a reinterpretation of $\vdash_{\mathcal{S}}$ incompatible with the present approach.

As usual, the rules of $\mathcal{S}$ are *backward deterministic*, and the **Subformula Property** continues to hold.

The systems $\mathcal{N}$ and $\mathcal{S}$ are obviously equivalent in presence of the Cut Rule. In particular, the translation from $\mathcal{N}$ to $\mathcal{S}$ requires proving the rules of $\mathcal{N}$ as derived rules of $\mathcal{S}$: as usual, $\vee$-introduction of $\mathcal{N}$ corresponds to $\vee$-introduction on the right, and $\vee$-elimination can be dealt with by $\vee$- introduction on the left and Cut, as follows:

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta \qquad \dfrac{\Gamma, \alpha \vdash_{\mathcal{S}} \gamma \quad \Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \vee \beta \vdash_{\mathcal{S}} \gamma}}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \gamma.}$$

Also the translation from $\mathcal{S}$ to $\mathcal{N}$ poses no problem: $\vee$-introduction on the right corresponds to $\vee$-introduction of $\mathcal{N}$, and $\vee$-introduction on the left can be dealt with as follows: given $\Gamma, \alpha \vdash_{\mathcal{N}} \gamma$ and $\Gamma, \beta \vdash_{\mathcal{N}} \gamma$, we can use them in an application of $\vee$-elimination from $\alpha \vee \beta$ by discharging, respectively, α and β, and this produces a proof of γ from the assumptions Γ and $\alpha \vee \beta$.

To get the full equivalence between the two systems, we need to extend 1.3.3 and 4.1.6.

Theorem 17.1.6 Cut Elimination (Gentzen [1935]) *For any Γ and β:*

$$\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \beta \;\Rightarrow\; \Gamma \vdash_{\mathcal{S}} \beta.$$

Proof. The Cut Elimination procedure of p. 22 (to which we refer in the following) can be extended to take care of the new connective $\vee$.

A cut was called *inductive* when the formula which is cut has just been introduced on both sides, for example:

$$\frac{\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha}{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta} \qquad \dfrac{\Gamma, \alpha \vdash_{\mathcal{S}} \gamma \quad \Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \vee \beta \vdash_{\mathcal{S}} \gamma}}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \gamma.}$$

Such a cut can be eliminated as follows, by substituting it with one cut on the formula α (of lower degree):

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \alpha \vdash_{\mathcal{S}} \gamma}{\Gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \gamma.}$$

A cut was called *interlocutory* when the formula which is cut has been introduced at steps preceeding the last ones (on the appropriate sides). In this case we simply move the cut upwards, until it can be eliminated as above. For example, a cut like

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \dfrac{\Gamma, \alpha, \gamma \vdash_{\mathcal{S}} \beta \qquad \Gamma, \alpha, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \alpha, \gamma \vee \delta \vdash_{\mathcal{S}} \beta}}{\Gamma, \gamma \vee \delta \vdash_{\mathcal{S}+\mathrm{Cut}} \beta}$$

can be replaced by *two* as

$$\frac{\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \alpha, \gamma \vdash_{\mathcal{S}} \beta}{\Gamma, \gamma \vdash_{\mathcal{S}+\mathrm{Cut}} \beta} \qquad \dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \alpha, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \delta \vdash_{\mathcal{S}+\mathrm{Cut}} \beta}}{\Gamma, \gamma \vee \delta \vdash_{\mathcal{S}} \beta.}$$

The remaining cases are similar. □

The Cut Elimination Theorem fills the remaining gap in the proof of the following result.

Corollary 17.1.7 Equivalence of the Sequent System and Natural Deduction (Gentzen [1935]) *For any* Γ *and* β*:*

$$\Gamma \vdash_{\mathcal{S}} \beta \Leftrightarrow \Gamma \vdash_{\mathcal{N}} \beta.$$

A different proof of Cut Elimination comes from the following extension of 1.3.6 and 4.1.8.

Proposition 17.1.8 (Prawitz [1965]) *There are canonical translations of cut-free proofs in* $\mathcal{S}$ *to normal proofs in* $\mathcal{N}$*, and conversely.*

Proof. The translation from $\mathcal{S}$ to $\mathcal{N}$ given as half of the proof of the equivalence of the two systems already shows that (cut-free) proofs in $\mathcal{S}$ correspond to normal proofs in $\mathcal{N}$.

For the converse, we only have to supplement 1.3.6 and 4.1.8. The case of the *introduction rules* of $\mathcal{N}$ can be dealt with by induction and the corresponding introduction rules on the right. For example, if $\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta$ has been obtained by $\vee$-introduction from a normal proof $\Gamma \vdash_{\mathcal{N}} \alpha$, then $\Gamma \vdash_{\mathcal{S}} \alpha$ by the induction hypothesis, and thus $\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta$ by $\vee$-introduction on the right.

The case of the *elimination rules* of $\mathcal{N}$ is the crucial one, and requires more ingenuity (since the natural translation uses the Cut Rule). The general schema has been shown in 1.3.6: if $\Gamma \vdash_{\mathcal{N}} \beta$ has been obtained by an elimination rule, the latter is either a $\rightarrow$-elimination, a $\wedge$-elimination, or a $\vee$-elimination. The three cases are treated similarly, by climbing up in the given proof until an assumption is reached that is *eliminated* in the first step of the given proof below it: this is possible because the given proof is normal, and β has been obtained by an elimination rule. Such an assumption must be of one of the following three forms: $\gamma \rightarrow \delta$, $\gamma \wedge \delta$, and $\gamma \vee \delta$. The first two cases have already been dealt with in 1.3.6 and 4.1.8, and we consider now the last one. The given proof is then, for example, of the form:

$$\Gamma,\ \dfrac{\gamma \vee \delta \qquad \begin{array}[b]{c}\Gamma, [\gamma]\\ \mathcal{D}^{\gamma}_{\alpha}\\ \alpha\end{array} \qquad \begin{array}[b]{c}\Gamma, [\delta]\\ \mathcal{D}^{\delta}_{\alpha}\\ \alpha\end{array}}{\begin{array}[t]{c}\alpha\\ \mathcal{D}_{\beta}\\ \beta.\end{array}} \tag{17.1}$$

We can then apply the induction hypothesis to the normal proofs

$$\begin{array}{cc}
 & \Gamma, \gamma\\
 & \mathcal{D}^{\gamma}_{\alpha}\\
\Gamma, & \alpha\\
 & \mathcal{D}_{\beta}\\
 & \beta
\end{array}
\qquad \text{and} \qquad
\begin{array}{cc}
 & \Gamma, \delta\\
 & \mathcal{D}^{\delta}_{\alpha}\\
\Gamma, & \alpha\\
 & \mathcal{D}_{\beta}\\
 & \beta
\end{array} \tag{17.2}$$

and get cut-free proofs $\Gamma, \gamma \vdash_{\mathcal{S}} \beta$ and $\Gamma, \delta \vdash_{\mathcal{S}} \beta$. By $\vee$-introduction on the left,

$$\frac{\Gamma, \gamma \vdash_{\mathcal{S}} \beta \qquad \Gamma, \delta \vdash_{\mathcal{S}} \beta}{\Gamma, \gamma \vee \delta \vdash_{\mathcal{S}} \beta.}$$

But $\gamma \vee \delta$ (being an assumption) is already in Γ, and thus the conclusion is equivalent to $\Gamma \vdash_{\mathcal{S}} \beta$. $\quad\square$

Notice that it was crucial to consider the extended notion of normal proofs, defined as absence of *delayed* maxima. Otherwise, if one of $\mathcal{D}^{\gamma}_{\alpha}$ and $\mathcal{D}^{\delta}_{\alpha}$ ends with an introduction and $\mathcal{D}_{\beta}$ starts with an elimination, α might become a simple maximum in one of 17.2, without being a simple maximum in 17.1. Of course, this cannot happen if α is not a delayed maximum, i.e. if the original proof is normal in the strong sense.

As usual, *the translation from $\mathcal{S}$ to $\mathcal{N}$ is not one-one.*

Kripke models

The notions of *intuitionistic possible world* and *intuitionistic logical consequence* do not refer to connectives, and can thus be retained in their original forms 2.2.1 and 2.2.4. What needs to be supplemented is the definition of forcing 2.2.2 and 4.2.1. This is where the notions of *Beth* and *Kripke forcing*, that coincided for the fragment consisting only of → and ∧, diverge. We first deal with the latter.

Definition 17.1.9 Kripke Forcing (Cohen [1963], Kripke [1963]) *For a given possible world $\mathcal{A}$, the relation $\Vdash_{\mathcal{A}}$ defined in 2.2.2 and 4.2.1 is extended to disjunction as follows:*

$$\sigma \Vdash_{\mathcal{A}} \alpha \vee \beta \quad \Leftrightarrow \quad \sigma \Vdash_{\mathcal{A}} \alpha \ \ \text{or} \ \ \sigma \Vdash_{\mathcal{A}} \beta.$$

The next result shows that the extension of forcing to ∨ captures the intended meaning of disjunction.

Theorem 17.1.10 Kripke Soundness and Completeness (Kripke [1963]) *For any Γ and α:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \ \Leftrightarrow \ \Gamma \models_i \alpha.$$

Proof. For the Soundness direction, we supplement the proof of 2.2.5 by the cases dealing with disjunction.

- If $\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta$ is obtained from, say, $\Gamma \vdash_{\mathcal{N}} \alpha$ by ∨-*introduction*, then $\Gamma \models_i \alpha$ by the induction hypothesis. Let σ be any state that forces all formulas of Γ in some world $\mathcal{A}$: then σ forces α, and hence it forces $\alpha \vee \beta$ by definition of forcing. Since σ and $\mathcal{A}$ are arbitrary, $\Gamma \models_i \alpha \vee \beta$. Similarly for the other ∨-introduction rule.

- If $\Gamma \vdash_{\mathcal{N}} \gamma$ is obtained from

  $$\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta \qquad \Gamma, \alpha \vdash_{\mathcal{N}} \gamma \qquad \Gamma, \beta \vdash_{\mathcal{N}} \gamma$$

  by ∨-*elimination*, then

  $$\Gamma \models_i \alpha \vee \beta \qquad \Gamma, \alpha \models_i \gamma \qquad \Gamma, \beta \models_i \gamma$$

  by the induction hypothesis. Let σ be any state that forces all formulas of Γ in some world $\mathcal{A}$. By the first induction hypothesis σ forces $\alpha \vee \beta$, and hence it forces α or β by definition of forcing. Then σ forces γ by the second or third induction hypothesis, respectively. Since σ and $\mathcal{A}$ are arbitrary, $\Gamma \models_i \gamma$.

For the Completeness direction, we supplement the proof of 2.2.6 (to which we refer) by the case of disjunction. We have to prove that

$$\Theta \Vdash_{\mathcal{A}} \gamma \vee \delta \ \Leftrightarrow \ \gamma \vee \delta \in \Theta,$$

where Θ is a set of formulas closed under $\vdash_{\mathcal{N}}$.

- Suppose $\Theta \Vdash_{\mathcal{A}} \gamma \vee \delta$. Then $\Theta \Vdash_{\mathcal{A}} \gamma$ or $\Theta \Vdash_{\mathcal{A}} \delta$, and by the induction hypothesis $\gamma \in \Theta$ or $\delta \in \Theta$. We want $\gamma \vee \delta \in \Theta$. Suppose $\gamma \vee \delta \notin \Theta$. By closure under $\vdash_{\mathcal{N}}$, $\Theta \nvdash_{\mathcal{N}} \gamma \vee \delta$. By $\vee$-*introduction*, $\Theta \nvdash_{\mathcal{N}} \gamma$ and $\Theta \nvdash_{\mathcal{N}} \delta$, contradiction (because Θ contains one of γ and δ).

- Suppose $\gamma \vee \delta \in \Theta$. We want $\Theta \Vdash_{\mathcal{A}} \gamma \vee \delta$, i.e. $\Theta \Vdash_{\mathcal{A}} \gamma$ or $\Theta \Vdash_{\mathcal{A}} \delta$ or, by the induction hypothesis, $\gamma \in \Theta$ or $\delta \in \Theta$. Thus we need to show that if $\gamma \vee \delta$ is in Θ, then so is one of γ and δ: this is not quite implied by $\vee$-elimination, but can easily be obtained by restricting attention to *saturated sets* of formulas, defined precisely by the condition that, for every γ and δ,

$$\gamma \vee \delta \in \Theta \;\Rightarrow\; \gamma \in \Theta \text{ or } \delta \in \Theta.$$

It follows that, if $\mathcal{A}$ is the world defined as follows:

$$\mathcal{A} = \langle \mathcal{F}, \subseteq, \{\mathcal{A}_\Theta\}_{\Theta \in \mathcal{F}} \rangle,$$

where:

1. $\mathcal{F}$ is the set of all saturated sets of formulas Θ closed under $\vdash_{\mathcal{N}}$
2. $\subseteq$ is the usual inclusion relation
3. $\mathcal{A}_\Theta$ is the set of formulas in Θ consisting only of a propositional letter,

then, for any formula α,

$$\Theta \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \alpha \in \Theta.$$

It remains to finish the proof as usual. We only need to show that *if* $\Gamma \nvdash_{\mathcal{N}} \alpha$, *then there is a saturated set* Θ *closed under* $\vdash_{\mathcal{N}}$ *such that* $\Theta \supseteq \Gamma$ *and* $\alpha \notin \Theta$ (since then Θ forces every formula in Γ but not α). This fact, that would be trivial if saturation were not required, as in 2.2.6 and 4.2.2 (because then Θ can simply be taken as the closure of Γ under $\vdash_{\mathcal{N}}$), requires now a proof. We provide two different ones, based on different principles.

For simplicity, we let

$$cl_{\mathcal{N}}(\Theta) = \text{the closure of } \Theta \text{ under } \vdash_{\mathcal{N}}.$$

1. *saturation*
   Let $\{\gamma_n \vee \delta_n\}_{n \in \omega}$ be a list of all disjunctions, in which each one occurs with infinitely many repetitions. We define Θ by approximations, as follows:

$$\begin{array}{lcl} \Theta_0 & = & cl_{\mathcal{N}}(\Gamma) \\ \Theta_{n+1} & = & \begin{cases} \Theta_n & \text{if } \Theta_n \nvdash_{\mathcal{N}} \gamma_n \vee \delta_n \\ cl_{\mathcal{N}}(\Theta_n \cup \{\gamma_n\}) & \text{if } \Theta_n \vdash_{\mathcal{N}} \gamma_n \vee \delta_n \text{ and } \Theta_n, \gamma_n \nvdash_{\mathcal{N}} \alpha \\ cl_{\mathcal{N}}(\Theta_n \cup \{\delta_n\}) & \text{otherwise} \end{cases} \\ \Theta & = & \bigcup_{n \in \omega} \Theta_n. \end{array}$$

Θ is closed under $\vdash_{\mathcal{N}}$ because if $\Theta \vdash_{\mathcal{N}} \beta$ then, since a deduction can only use finitely many premises, $\Theta_{n+1} \vdash_{\mathcal{N}} \beta$ for some n, and each Θ_{n+1} is closed under $\vdash_{\mathcal{N}}$ by definition.

Θ is saturated because if $\Theta \vdash_{\mathcal{N}} \gamma \vee \delta$ then, as above, $\Theta_m \vdash_{\mathcal{N}} \gamma \vee \delta$ for some m. Since each disjunction appears infinitely often in the given list, there is $n \geq m$ such that $\gamma \vee \delta = \gamma_n \vee \delta_n$, and thus $\Theta_n \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$ because $\Theta_n \supseteq \Theta_m$. By construction, one of $\gamma = \gamma_n$ and $\delta = \delta_n$ will then go into Θ_{n+1}, and hence in Θ.

$\Theta \not\vdash_{\mathcal{N}} \alpha$ since, by induction on n, $\Theta_n \not\vdash_{\mathcal{N}} \alpha$. For $n = 0$ this is simply the hypothesis on Γ. And for $n+1$ this is so by construction, since: either $\Theta_n \not\vdash_{\mathcal{N}} \gamma_n \vee \delta_n$ (and then there is nothing to prove, because $\Theta_{n+1} = \Theta_n$); or $\Theta_n \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$, and then it cannot be both

$$\Theta_n, \gamma_n \vdash_{\mathcal{N}} \alpha \qquad \text{and} \qquad \Theta_n, \delta_n \vdash_{\mathcal{N}} \alpha$$

otherwise, by $\vee$-*elimination* (that we had not used yet!),

$$\Theta_n, \gamma_n \vee \delta_n \vdash_{\mathcal{N}} \alpha$$

and thus $\Theta_n \vdash_{\mathcal{N}} \alpha$ (since $\Theta_n \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$), contradicting the induction hypothesis on n. By definition, then Θ_{n+1} is chosen in such a way that $\Theta_{n+1} \not\vdash_{\mathcal{N}} \alpha$ either.

Finally, $\Theta \supseteq \Gamma$ by definition of Θ_0.

2. *maximality*

Look at the collection of sets of formulas closed under $\vdash_{\mathcal{N}}$, including Γ, and not containing α. Let Θ be any set maximal (under inclusion) in this collection.

We show that Θ is saturated, as follows. If $\Theta \vdash_{\mathcal{N}} \gamma \vee \delta$, suppose

$$\Theta \not\vdash_{\mathcal{N}} \gamma \qquad \text{and} \qquad \Theta \not\vdash_{\mathcal{N}} \delta.$$

Then

$$\Theta, \gamma \vdash_{\mathcal{N}} \alpha \qquad \text{and} \qquad \Theta, \delta \vdash_{\mathcal{N}} \alpha$$

by maximality,

$$\Theta, \gamma \vee \delta \vdash_{\mathcal{N}} \alpha$$

by $\vee$-*elimination*, and

$$\Theta \vdash_{\mathcal{N}} \alpha$$

because $\gamma \vee \delta \in \Theta$ (from $\Theta \vdash_{\mathcal{N}} \gamma \vee \delta$, by closure under $\vdash_{\mathcal{N}}$), contradiction.

It remains to argue that such a maximal set Θ exists. We can either appeal to Zorn's Lemma, or build Θ directly as follows. Let $\{\beta_n\}_{n\in\omega}$ be an enumeration of all formulas. Let

$$\begin{array}{lcl} \Theta_0 & = & cl_{\mathcal{N}}\Gamma \\ \Theta_{n+1} & = & \left\{ \begin{array}{ll} cl_{\mathcal{N}}(\Theta_n \cup \{\beta_n\}) & \text{if } \Theta_n, \beta_n \nvdash_{\mathcal{N}} \alpha \\ \Theta_n & \text{otherwise} \end{array} \right. \\ \Theta & = & \bigcup_{n\in\omega} \Theta_n. \end{array}$$

Θ is closed under $\vdash_{\mathcal{N}}$ because if $\Theta \vdash_{\mathcal{N}} \beta$ then, since a deduction can only use finitely many premises, $\Theta_{n+1} \vdash_{\mathcal{N}} \beta$ for some n, and each Θ_{n+1} is closed under $\vdash_{\mathcal{N}}$ by definition.

Θ is maximal because if $\beta \notin \Theta$ and $\Theta \cup \{\beta\} \nvdash_{\mathcal{N}} \alpha$, there is n such that $\beta = \beta_n$. Then $\Theta_n \cup \{\beta_n\} \nvdash_{\mathcal{N}} \alpha$ (since $\Theta_n \subseteq \Theta$), and by construction $\beta = \beta_n \in \Theta_{n+1}$, i.e. $\beta \in \Theta$, contradiction.

$\Theta \nvdash_{\mathcal{N}} \alpha$ since, by induction on n, $\Theta_n \nvdash_{\mathcal{N}} \alpha$. For $n = 0$ this is simply the hypothesis on Γ. And for $n+1$ this is so by construction, since if $\Theta_n, \beta_n \nvdash_{\mathcal{N}} \alpha$ then α is not in the closure of $\Theta_n \cup \{\beta_n\}$, i.e. in Θ_{n+1}; and if $\Theta_n, \beta_n \vdash_{\mathcal{N}} \alpha$ then $\Theta_{n+1} = \Theta_n$, and by the induction hypothesis $\Theta_n \nvdash_{\mathcal{N}} \alpha$. □

In particular, we get a *semantical proof of the Cut Elimination Theorem*, as in 2.2.7.

Exercise 17.1.11 a) *The Countable Model Property continues to hold.* (Hint: see 2.2.9. Now not only closure under $\mathcal{N}$ has to be ensured, but also saturation.)

b) *The Finite Model Property continues to hold.* (Hint: see 2.2.11.)

Beth models

There is no problem in extending the proof of the Countable Model Property (2.2.9) to the case of disjunction, and we left this as an exercise above: basically, we consider as nodes all finitely generated, saturated extensions of Γ closed under $\vdash_{\mathcal{N}}$.

Toward an extension of the Constructive Model Property (2.2.10), there is also no problem in avoiding the consideration of closure under $\vdash_{\mathcal{N}}$. It remains to deal with the problem of avoiding the saturation condition, which is by itself an infinitary condition.

We were able to avoid closure under $\vdash_{\mathcal{N}}$ simply by forgetting about it, since closure was used only to collapse deducibility and membership (in other words, for reasons of elegance). Saturation has instead a more substantial use in dealing with disjunction, and it cannot simply be forgotten about. What we can do, however, is to ensure saturation only in the limit, i.e. for *branches* of the tree, as opposed to in the present, i.e. for *nodes*.

This can be achieved by mixing the usual construction with the saturation procedure of 17.1.10. Basically, we will make sure that whenever a disjunction $\gamma \vee \delta$ is deducible from a node T_σ, then one of γ and δ will eventually be in every branch extending T_σ (not necessarily the same one for every branch).

This obviously saturates the branches of the model but not the nodes, and thus we cannot hope to prove

$$\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \Gamma_\sigma \vdash_{\mathcal{N}} \alpha$$

in the case of disjunctions. After all, if a node forces $\gamma \vee \delta$, then it forces one of γ or δ, while we only ensure that if T_σ deduces $\gamma \vee \delta$, then every branch going through T_σ will eventually deduce one of γ or δ. Since we have no room for substantial improvements in the construction (we are trying to replace an infinitary step by finitary ones, and this can be achieved only in the limit), we have to play somewhere else.

We do this by *modifying the notion of forcing*, and adapting it to what the construction achieves. Precisely, we require that a node forces $\gamma \vee \delta$ precisely when one of γ or δ will eventually be forced in the future of that node, for every possible future: in other words, when any branch going through the node eventually reaches a point in which one of γ or δ is forced. As we might expect, since forcing eventually relies on atomic statements, a similar modification will be needed for forcing of propositional letters.

Technically, we can express the notion of a future of a node σ by considering any *bar* B_σ, i.e. any collection of states extending σ such that any (maximal) branch going through σ intersects B_σ, in the sense of going through one of its elements.

Definition 17.1.12 Beth Forcing (Beth [1956]) *For a given possible world $\mathcal{A}$, the relation $\Vdash_{\mathcal{A}}$ defined in 17.1.9 is modified as follows:*

$$\begin{array}{lcl} \sigma \Vdash_{\mathcal{A}} p & \Leftrightarrow & (\exists B_\sigma)(\forall \tau \in B_\sigma)(p \in \mathcal{A}_\tau) \\ \sigma \Vdash_{\mathcal{A}} \alpha \vee \beta & \Leftrightarrow & (\exists B_\sigma)(\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \alpha \;\; or \;\; \tau \Vdash_{\mathcal{A}} \beta). \end{array}$$

The notion of Kripke forcing is recovered as a special case, when $B_\sigma = \{\sigma\}$ for every σ.

We will consider only intuitionistic worlds whose underlying partial order is a binary tree, since they are the ones used in the previous proofs of the Constructive Model Property. The restriction will have certain advantages, in particular the fact that we will be able to restrict attention to bars of finite height (this is needed in the proof of 17.1.14).[4]

[4] By adopting this restricted notion of intuitionistic world, we can *prove* the properties of bars needed in the following proofs. Alternately, we could leave the notion of intuitionistic world unchanged, and take the needed properties of bars as part of the *definition* of a bar. As shown by the proof of 17.1.14, these properties are the following:

1. $\{\sigma\}$ is a bar for σ

Notice that *Beth forcing is monotone*: the atomic case holds by monotonicity of knowledge (which still holds by definition); for disjunction, if $\sigma \Vdash_{\mathcal{A}} \alpha \vee \beta$ then there is a bar B_σ such that every state in it forces either α or β, and if $\tau \sqsupseteq \sigma$ then a bar for τ can be obtained by taking all states in B_σ that extend τ; the other cases are as for Kripke forcing.

We talk of **Beth models** or **Kripke models** according to the notion of forcing we use, on top of the associated notion of intuitionistic world. In particular, there are now two notions of intuitionistic logical consequence (and, as a limit case, of intuitionistic validity):

- $\boldsymbol{\Gamma \models_i \alpha}$ means that α is Kripke forced in every world, at every state in which all formulas of Γ are Kripke forced
- $\boldsymbol{\Gamma \models_{ib} \alpha}$ means that α is Beth forced in every world, at every state in which all formulas of Γ are Beth forced.

The next result spells out the connections between the two kinds of forcing.

Proposition 17.1.13 *For every Kripke model $\mathcal{A}$ there exists a Beth model $\mathcal{B}$ such that the same formulas are (globally) forced on $\mathcal{A}$ and $\mathcal{B}$, but the converse does not hold.*

Proof. Given a Kripke model

$$\mathcal{A} = \langle P_{\mathcal{A}}, \sqsubseteq_{\mathcal{A}}, \{\mathcal{A}_\sigma\}_{\sigma \in P_{\mathcal{A}}} \rangle,$$

we define a Beth model

$$\mathcal{B} = \langle P_{\mathcal{B}}, \sqsubseteq_{\mathcal{B}}, \{\mathcal{B}_\sigma\}_{\sigma \in P_{\mathcal{B}}} \rangle$$

as follows:

- the elements of $P_{\mathcal{B}}$ are the finite sequences $\langle \sigma_0, \ldots, \sigma_n \rangle$ of elements of $P_{\mathcal{A}}$, such that $\sigma_0 \sqsubseteq_{\mathcal{A}} \cdots \sqsubseteq_{\mathcal{A}} \sigma_n$
- $\sqsubseteq_{\mathcal{B}}$ is the order of sequences by initial segments
- $\mathcal{B}_{\langle \sigma_0, \ldots, \sigma_n \rangle} = \mathcal{A}_{\sigma_n}$.

We now show that

$$\langle \sigma_0, \ldots, \sigma_n \rangle \Vdash_{\mathcal{B}} \alpha \iff \sigma_n \Vdash_{\mathcal{A}} \alpha$$

by induction on α, the only interesting cases being the left-to-right directions when α is atomic or a disjunction.

2. if B_σ is a bar for σ and, for every $\tau \in B_\sigma$, B_τ is a bar for τ, then $\bigcup_{\tau \in B_\sigma} B_\tau$ is still a bar for σ (the bar of a bar is a bar)
3. if B_σ is a bar for σ and $\tau \sqsupseteq \sigma$, then the intersection of B_σ with the set of all extensions of τ is a bar for τ.

Incidentally, these axioms define the socalled *Grothendieck topology*.

- *propositional letters*
  If $\langle \sigma_0, \ldots, \sigma_n \rangle \Vdash_{\mathcal{B}} p$ then, by definition of Beth forcing, every branch through $\langle \sigma_0, \ldots, \sigma_n \rangle$ will eventually reach a point τ such that $p \in \mathcal{B}_\tau$. Consider the branch $\langle \sigma_0, \ldots, \sigma_n, \sigma_n, \sigma_n, \ldots \rangle$. Then, for the right number of repetions,

  $$p \in \mathcal{B}_{\langle \sigma_0, \ldots, \sigma_n, \ldots, \sigma_n \rangle} = \mathcal{A}_{\sigma_n},$$

  so that $\sigma_n \Vdash_{\mathcal{A}} p$ by definition of Kripke forcing.

- *disjunction*
  If $\langle \sigma_0, \ldots, \sigma_n \rangle \Vdash_{\mathcal{B}} p$ then, by definition of Beth forcing, every branch through $\langle \sigma_0, \ldots, \sigma_n \rangle$ will eventually reach a point τ such that $\tau \Vdash_{\mathcal{B}} \alpha$ or $\tau \Vdash_{\mathcal{B}} \beta$. Consider the branch $\langle \sigma_0, \ldots, \sigma_n, \sigma_n, \sigma_n, \ldots \rangle$. Then, for the right number of repetions,

  $$\langle \sigma_0, \ldots, \sigma_n, \ldots, \sigma_n \rangle \Vdash_{\mathcal{B}} \alpha \text{ or } \langle \sigma_0, \ldots, \sigma_n, \ldots, \sigma_n \rangle \Vdash_{\mathcal{B}} \beta.$$

  By induction hypothesis $\sigma_n \Vdash_{\mathcal{A}} \alpha$ or $\sigma_n \Vdash_{\mathcal{A}} \beta$, and hence $\sigma_n \Vdash_{\mathcal{A}} \alpha \vee \beta$ by definition of Kripke forcing.

For the counterexample to the opposite implication, consider the Beth model consisting of the nodes $\emptyset, \langle 0 \rangle$ and $\langle 1 \rangle$, and such that $\mathcal{B} = \emptyset$, $\mathcal{B}_{\langle 0 \rangle} = \{p\}$ and $\mathcal{B}_{\langle 1 \rangle} = \{q\}$. By definition of Beth forcing, $p \vee q$ is forced in $\mathcal{B}$, although neither p nor q are. But by definition of Kripke forcing, this cannot happen in any Kripke model. □

Notice that a finite Kripke model becomes an infinite Beth model in the above translation, as every node is repeated any finite number of times. This is necessary, because the Finite Model Property holds for Kripke models but not for Beth models, since *any finite Beth model forces* $\alpha \vee \neg\alpha$. Indeed, for each terminal node σ of a model, either σ forces α or no extension of it does, and hence σ forces $\neg\alpha$. And for Beth models, although not for Kripke models, this is enough to force $\alpha \vee \neg\alpha$ globally.

We now prove soundness with respect to Beth models, which in view of the previous translation is stronger than that with respect to Kripke models.

Theorem 17.1.14 Beth Soundness (Beth [1956]) *For any* Γ *and* α*:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \;\Rightarrow\; \Gamma \models_{ib} \alpha.$$

Proof. We only have to deal with disjunction, since the other connectives are dealt with as in Kripke forcing.

- If $\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta$ is obtained from, say, $\Gamma \vdash_{\mathcal{N}} \alpha$ by $\vee$-*introduction*, then $\Gamma \models_{ib} \alpha$ by the induction hypothesis. Let σ be any state that Beth forces all formulas of Γ in some world $\mathcal{A}$: then σ Beth forces α, and hence it forces $\alpha \vee \beta$ by definition of Beth forcing (taking $B_\sigma = \{\sigma\}$, which is obviously a bar for σ^5). Since σ and $\mathcal{A}$ are arbitrary, $\Gamma \models_{ib} \alpha \vee \beta$. Similarly for the other $\vee$-introduction rule.

- If $\Gamma \vdash_{\mathcal{N}} \gamma$ is obtained from

$$\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta \qquad \Gamma, \alpha \vdash_{\mathcal{N}} \gamma \qquad \Gamma, \beta \vdash_{\mathcal{N}} \gamma$$

by $\vee$-*elimination*, then

$$\Gamma \models_{ib} \alpha \vee \beta \qquad \Gamma, \alpha \models_{ib} \gamma \qquad \Gamma, \beta \models_{ib} \gamma$$

by the induction hypothesis. Let σ be any state that Beth forces all formulas of Γ in some world $\mathcal{A}$. By the first induction hypothesis, σ Beth forces $\alpha \vee \beta$, and hence there is B_σ such that

$$(\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \alpha \ \text{ or } \ \tau \Vdash_{\mathcal{A}} \beta).$$

By the second and third induction hypotheses,

$$(\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \gamma).$$

It is then enough to prove that if a formula is Beth forced at every node of a bar of σ, then it is Beth forced at σ itself.

To conclude the proof, it thus remains to show that

$$(\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \gamma) \ \Rightarrow \ (\sigma \Vdash_{\mathcal{A}} \gamma).$$

We proceed by induction on γ.

1. *propositional letters*
If $\gamma = p$ and $\tau \Vdash_{\mathcal{A}} p$ then, by definition of Beth forcing, there is a bar B_τ such that

$$(\forall \nu \in B_\tau)(\nu \Vdash_{\mathcal{A}} p).$$

If we let

$$B^*_\sigma = \bigcup_{\tau \in B_\sigma} B_\tau$$

then B^*_σ is still a bar for σ^6 because every branch through σ goes through some $\tau \in B_\sigma$, every branch through τ goes through some $\nu \in B_\tau$, and hence every branch through σ goes some through some $\nu \in B^*_\sigma$. Moreover,

$$(\forall \nu \in B^*_\sigma)(\nu \Vdash_{\mathcal{A}} p),$$

and hence $\sigma \Vdash_{\mathcal{A}} p$ by definition of Beth forcing.

[5] Cf. axiom 1 in note 4.

[6] Cf. axiom 2 of note 4.

2. *implication*

   If $\gamma = \gamma_1 \to \gamma_2$, then
   $$(\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \gamma_1 \to \gamma_2)$$
   by hypothesis, and
   $$(\forall \tau \in B_\sigma)(\forall \nu \sqsupseteq \tau)(\nu \Vdash_{\mathcal{A}} \gamma_1 \ \Rightarrow \nu \Vdash_{\mathcal{A}} \gamma_2)$$
   by definition of forcing. We want to show that $\sigma \Vdash_{\mathcal{A}} \gamma_1 \to \gamma_2$, i.e.
   $$(\forall \tau \sqsupseteq \sigma)(\tau \Vdash_{\mathcal{A}} \gamma_1 \ \Rightarrow \ \tau \Vdash_{\mathcal{A}} \gamma_2).$$
   Given any $\tau \sqsupseteq \sigma$, consider the set B_τ of all states ν in B_σ that are extensions of τ: by definition of a bar, this is still a bar for τ.[7] If $\tau \Vdash_{\mathcal{A}} \gamma_1$, then $\nu \Vdash_{\mathcal{A}} \gamma_1$ for all $\nu \in B_\tau$ by monotonicity of forcing, and thus $\nu \Vdash_{\mathcal{A}} \gamma_2$ by hypothesis (since $\nu \sqsupseteq \nu \in B_\tau$). Thus
   $$(\forall \nu \in B_\tau)(\nu \Vdash_{\mathcal{A}} \gamma_2),$$
   and then $\tau \Vdash_{\mathcal{A}} \gamma_2$ by the induction hypothesis.

3. *conjunction*

   If $\gamma = \gamma_1 \wedge \gamma_2$, then
   $$\begin{aligned} & (\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \gamma_1 \wedge \gamma_2) \\ \Leftrightarrow\quad & (\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \gamma_1 \text{ and } \tau \Vdash_{\mathcal{A}} \gamma_2) \\ \Leftrightarrow\quad & (\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \gamma_1) \text{ and } (\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \gamma_2) \\ \Leftrightarrow\quad & (\sigma \Vdash_{\mathcal{A}} \gamma_1) \text{ and } (\sigma \Vdash_{\mathcal{A}} \gamma_2) \\ \Leftrightarrow\quad & \sigma \Vdash_{\mathcal{A}} \gamma_1 \wedge \gamma_2 \end{aligned}$$
   by definition of forcing and induction hypothesis.

4. *disjunction*

   If $\gamma = \gamma_1 \vee \gamma_2$ and $\tau \Vdash_{\mathcal{A}} \gamma_1 \vee \gamma_2$ then, by definition of Beth forcing, there is a bar B_τ such that
   $$(\forall \nu \in B_\tau)(\nu \Vdash_{\mathcal{A}} \gamma_1 \text{ or } \nu \Vdash_{\mathcal{A}} \gamma_2).$$
   If we let
   $$B^*_\sigma = \bigcup_{\tau \in B_\sigma} B_\tau$$
   then, as in case 1, B^*_σ is still a bar for σ.[8] Moreover,
   $$(\forall \nu \in B^*_\sigma)(\nu \Vdash_{\mathcal{A}} \gamma_1 \text{ or } \nu \Vdash_{\mathcal{A}} \gamma_2),$$
   and hence $\sigma \Vdash_{\mathcal{A}} \gamma_1 \vee \gamma_2$ by definition of Beth forcing. □

[7] Cf. axiom 3 of note 4.

[8] Cf. axiom 2 of note 4.

Because of the previous translation, completeness with respect to Beth models is weaker than that with respect to Kripke models. The next result is thus a consequence of 17.1.10, but the direct proof given below produces the corollary which constitutes the original motivation for the introduction of Beth models.

Theorem 17.1.15 Beth Completeness (Beth [1956]) *For any* Γ *and* α*:*

$$\Gamma \models_{ib} \alpha \ \Rightarrow\ \Gamma \vdash_{\mathcal{N}} \alpha.$$

Proof. We alternate steps to get all finitely generated extensions of Γ, to steps to get saturation in the limit. Precisely, let $\{\alpha_n\}_{n\in\omega}$ be an enumeration of all formulas (in the language $\to$, $\wedge$, and $\vee$), and $\{\gamma_n \vee \delta_n\}_{n\in\omega}$ be an enumeration of all disjunctions (in the same language), both with infinitely many repetitions. We start with

$$\Gamma_\emptyset = \Gamma.$$

If $|\sigma| = 2n$, then

$$\Gamma_{\sigma*\langle 0\rangle} = \Gamma_\sigma \qquad \text{and} \qquad \Gamma_{\sigma*\langle 1\rangle} = \Gamma_\sigma \cup \{\alpha_n\}.$$

If $|\sigma| = 2n+1$, then: if $\Gamma_\sigma \not\vdash_{\mathcal{N}} \gamma_n \vee \delta_n$,

$$\Gamma_{\sigma*\langle 0\rangle} = \Gamma_{\sigma*\langle 1\rangle} = \Gamma_\sigma;$$

and if $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$,

$$\Gamma_{\sigma*\langle 0\rangle} = \Gamma_\sigma \cup \{\gamma_n\} \qquad \text{and } \Gamma_{\sigma*\langle 1\rangle} = \Gamma_\sigma \cup \{\delta_n\}.$$

The world $\mathcal{A}$ is defined as follows:

$$\mathcal{A} = \langle \{\Gamma_\sigma\}_{\sigma\in\mathcal{S}}, \subseteq, \{\mathcal{A}_{\Gamma_\sigma}\}_{\sigma\in\mathcal{S}}\rangle,$$

where:

1. $\mathcal{S}$ is the set of all sequences of 0's and 1's
2. $\subseteq$ is the usual set-theoretical inclusion relation
3. $\mathcal{A}_{\Gamma_\sigma}$ is the set of propositional letters deducible from Γ_σ, i.e.

$$\mathcal{A}_{\Gamma_\sigma} = \{p : \Gamma_\sigma \vdash_{\mathcal{N}} p\}.$$

For simplicity of notations, we identify bars for Γ_σ with bars for σ that generate them, and we write B_σ for B_{Γ_σ}.

As usual, we want to prove that, for any formula α and string σ,

$$\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha \ \Leftrightarrow\ \Gamma_\sigma \vdash_{\mathcal{N}} \alpha.$$

We proceed by induction on α

1. *propositional letters*
If $\Gamma_\sigma \vdash_{\mathcal{N}} p$, then $p \in \mathcal{A}_{\Gamma_\sigma}$ by definition. If we let $B_\sigma = \{\sigma\}$, then
$$(\forall \tau \in B_\sigma)(p \in \mathcal{A}_{\Gamma_\tau}),$$
i.e. $\Gamma_\sigma \Vdash_{\mathcal{A}} p$.
Conversely, if $\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha$ then, by definition of Beth forcing,
$$(\exists B_\sigma)(\forall \tau \in B_\sigma)(p \in \mathcal{A}_{\Gamma_\tau}).$$
By definition of $\mathcal{A}_{\Gamma_\tau}$,
$$(\exists B_\sigma)(\forall \tau \in B_\sigma)(\Gamma_\tau \vdash_{\mathcal{N}} p).$$
It is thus enough to show that if a formula is provable at Γ_τ for every τ in a bar of σ, then it is provable at Γ_σ itself. This we will do at the end.

2. *conjunction*
As in 4.2.3.b for Kripke forcing.

3. *disjunction*
If $\alpha = \alpha_1 \vee \alpha_2$ and $\Gamma_\sigma \vdash_{\mathcal{N}} \alpha_1 \vee \alpha_2$, then consider an n so that
$$\gamma_n \vee \delta_n = \alpha_1 \vee \alpha_2 \qquad \text{and} \qquad 2n+1 \geq |\sigma|,$$
which exists because $\alpha_1 \vee \alpha_2$ appears in the list $\{\gamma_n \vee \delta_n\}_{n\in\omega}$ infinitely often. Then, for any $\tau \sqsupseteq \sigma$ of length $2n+1$, $T_\tau \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$ (because T_τ extends T_σ, and $T_\sigma \vdash_{\mathcal{N}} \alpha_1 \vee \alpha_2$). By construction,
$$T_{\tau * \langle 0\rangle} = T_\tau \cup \{\alpha_1\} \qquad \text{and} \qquad T_{\tau * \langle 1\rangle} = T_\tau \cup \{\alpha_2\}.$$
In particular,
$$T_{\tau * \langle 0\rangle} \vdash_{\mathcal{N}} \alpha_1 \qquad \text{and} \qquad T_{\tau * \langle 1\rangle} \vdash_{\mathcal{N}} \alpha_2.$$
By the induction hypothesis,
$$T_{\tau * \langle 0\rangle} \Vdash_{\mathcal{A}} \alpha_1 \qquad \text{and} \qquad T_{\tau * \langle 1\rangle} \Vdash_{\mathcal{A}} \alpha_2.$$
If we let B_σ be the bar of σ consisting of all strings ν extending σ and of length $2n+2$ (i.e. $\nu = \tau * \langle 0\rangle$ or $\nu = \tau * \langle 1\rangle$ for τ as above), we then have
$$(\forall \nu \in B_\sigma)(\Gamma_\nu \Vdash_{\mathcal{A}} \alpha_1 \text{ or } \Gamma_\nu \Vdash_{\mathcal{A}} \alpha_2),$$
and thus
$$\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha_1 \vee \alpha_2$$

by definition of Beth forcing.

Conversely, if $\Gamma_\sigma \Vdash_{\mathcal{A}} \alpha_1 \vee \alpha_2$ then, by definition of Beth forcing,

$$(\exists B_\sigma)(\forall \tau \in B_\sigma)(\Gamma_\tau \Vdash_{\mathcal{A}} \alpha_1 \text{ or } \Gamma_\tau \Vdash_{\mathcal{A}} \alpha_2).$$

By $\vee$-*introduction*,

$$(\exists B_\sigma)(\forall \tau \in B_\sigma)(\Gamma_\tau \Vdash_{\mathcal{A}} \alpha_1 \vee \alpha_2).$$

As in part 1, it is thus enough to show that if a formula is provable at Γ_τ for every τ in a bar of σ, then it is provable at Γ_σ itself. This we do below.

4. *implication*
   As in 2.2.10 for Kripke forcing, with a little modification due to the fact that, given $\Gamma_\sigma, \gamma \not\vdash \delta$, in general we do not have $\tau \sqsupseteq \sigma$ such that $\Gamma_\tau = \Gamma_\sigma \cup \{\gamma\}$, but only $\Gamma_\tau \supseteq \Gamma_\sigma \cup \{\gamma\}$, because at odd stages we might add disjunts when needed. But we do this only when their disjunction is already provable, so it is enough to note that

   $$\text{if } \Gamma, \alpha \vee \beta \not\vdash \delta \text{ then } \Gamma, \alpha \vee \beta, \alpha \not\vdash \delta \text{ or } \Gamma, \alpha \vee \beta, \beta \not\vdash \delta.$$

   Otherwise $\Gamma, \alpha \vee \beta \vdash \alpha \to \delta$ and $\Gamma, \alpha \vee \beta \vdash \beta \to \delta$, and so $\Gamma, \alpha \vee \beta \vdash \delta$ by $\vee$-elimination. So we can *choose* a branch such that $\tau \sqsupseteq \sigma$ and $\Gamma_\tau \supseteq \Gamma_\sigma \cup \{\gamma\}$ and $\Gamma_\tau \not\vdash \delta$. Then we can continue as in 2.2.10.

To conclude the proof, it remains to show that

$$(\forall \tau \in B_\sigma)(\Gamma_\tau \vdash_{\mathcal{N}} \gamma) \;\Rightarrow\; (\Gamma_\sigma \vdash_{\mathcal{N}} \gamma).$$

Notice that this property is the exact analogue for provability of the property proved at the end of 17.1.14 for forcing.

We first notice that, by monotonicity of forcing and the use of binary trees as intuitionistic worlds, it is enough to consider *horizontal bars*, i.e. bars of the form

$$B_\sigma^n = \{\tau : \tau \sqsupseteq \sigma \;\wedge\; |\tau| = |\sigma| + n\}.$$

We can then proceed by induction on the level of the horizontal bar, i.e. on $n \geq 1$. The inductive step is trivial, and we thus concentrate on $n = 1$, where $B_\sigma^1 = \{\sigma * \langle 0 \rangle, \sigma * \langle 1 \rangle\}$.

If $|\sigma| = 2n$, or $|\sigma| = 2n + 1$ and $\Gamma_\sigma \not\vdash_{\mathcal{N}} \gamma_n \vee \delta_n$, there is nothing to prove: at least one of $\Gamma_{\sigma * \langle 0 \rangle}$ and $\Gamma_{\sigma * \langle 1 \rangle}$ is equal to Γ_σ. Thus, if γ is deducible from both of the former, it is also deducible from the latter.

If $|\sigma| = 2n + 1$ and $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$, then

$$\Gamma_{\sigma * \langle 0 \rangle} = \Gamma_\sigma \cup \{\gamma_n\} \qquad \text{and } \Gamma_{\sigma * \langle 1 \rangle} = \Gamma_\sigma \cup \{\delta_n\}.$$

If γ is deducible from both of them, then it is also deducible from $\Gamma_\sigma \cup \{\gamma_n \vee \delta_n\}$ by $\vee$-*elimination*, and hence from Γ_σ alone because $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$. □

Corollary 17.1.16 Constructive Model Property. *For any Γ there is a constructively presented world $\mathcal{A}_\Gamma$ in which all formulas of Γ are Beth forced and such that, for every α, if $\Gamma \vdash_{\mathcal{N}} \alpha$ fails, then α is not Beth forced in $\mathcal{A}_\Gamma$.*

Proof. An examination of the proof just given shows that it is constructive. The main point to notice is that the steps to ensure saturation in the limit involve checking whether $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$, for $|\sigma| = 2n+1$: this can be done effectively, since the relation $\vdash_{\mathcal{N}}$ is decidable. □

Actually, the appeal to decidability of the $\vdash_{\mathcal{N}}$ relation in the proof of the corollary can be avoided by noticing that we don't really have to check whether $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$. It is enough that we check whether it does in $|\sigma|$ steps and, if not, that we continue to check also at later stages, i.e. for more and more steps. If $\Gamma_\sigma \vdash_{\mathcal{N}} \gamma_n \vee \delta_n$, eventually we will discover it and be able to put γ_n on one branch and δ_n on the other.

Intuitionistic tableaux

The notion of *provability by intuitionistic tableaux* does not refer to connectives, and can thus be retained in the original form 2.3.2. What needs to be supplemented is the definition of tableaux 2.3.1 and 4.2.4.

Definition 17.1.17 *An* **intuitionistic tableau** *is a tree with nodes consisting of signed forcing assertions of the form $T\sigma \Vdash \alpha$ or $F\sigma \Vdash \alpha$, and consistent with the formation rules of 2.3.1 and 4.2.4, as well as with the following:*

5. *If a node $T\sigma \Vdash \alpha \vee \beta$ is on the tree, then we can split any branch going through it by adding $T\sigma \Vdash \alpha$ in one direction and $T\sigma \Vdash \beta$ in the other. Graphically,*

$$\begin{array}{c} T\sigma \Vdash \alpha \vee \beta \\ \hline\hline T\sigma \Vdash \alpha \qquad T\sigma \Vdash \beta, \end{array}$$

   where the double line shows that the bottom nodes do not have to immediately follow the top one.

6. *If a node $F\sigma \Vdash \alpha \vee \beta$ is on the tree, then we can extend any branch going through it by adding $F\sigma \Vdash \alpha$ and $F\sigma \Vdash \beta$. Graphically,*

$$\begin{array}{c} F\sigma \Vdash \alpha \vee \beta \\ \hline\hline F\sigma \Vdash \alpha \\ \hline F\sigma \Vdash \beta. \end{array}$$

The next result shows that the extension of the tableaux rules to $\vee$ captures the intended meaning of validity.

Theorem 17.1.18 Soundness and Completeness for Tableaux (Nerode [1990]) *For any* Γ *and* α*:*

$$\Gamma \vdash_{\mathcal{T}} \alpha \;\Leftrightarrow\; \Gamma \models_i \alpha.$$

Proof. The proofs of 2.3.4 and 2.3.5 are easily supplemented by the cases dealing with disjunction, since the rules for the construction of tableaux were chosen to mirror the definition of Kripke forcing. □

Heyting $\sqcup\sqcap$-algebras

We extend the algebraic approach to the case of $\vee$, following the treatment of Chapter 5. Basically, we only have to add an operation $\sqcup_{\mathcal{A}}$ intended to model $\vee$. As usual, we will drop the subscript $\mathcal{A}$ when no confusion arises.

Definition 17.1.19 Canonical Interpretation. *Given a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, \sqcap, \sqcup, \Rightarrow \rangle$$

and an **environment** ρ *on it, the canonical interpretation* $[\![\ ]\!]_\rho$ *defined in 5.1.1 is extended to disjunction as follows:*

$$[\![\alpha]\!]_\rho = [\![\beta]\!]_\rho \sqcup [\![\gamma]\!]_\rho \qquad \textit{if } \alpha = \beta \vee \gamma.$$

The definition of an algebraic model refers only to $\sqsubseteq$ and $\sqcap$, and it is unchanged. We already know that the structure of $\mathcal{N}$ induces a structure of Heyting $\sqcap$-algebra on the equivalence classes of formulas under provable equivalence. We now consider the additional conditions imposed by the rules for $\vee$, by resuming the discussion started on p. 66, and continuing its enumeration.

5. $\vee$ *induces a least upper bound operation on the equivalence classes*
   The $\vee$-introduction rules show that $\vee$ induces an *upper bound*:

   $$\frac{\Gamma \vdash_{\mathcal{N}} \alpha}{\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta}$$

   says that anything less than α must be less than $\alpha \vee \beta$. Similarly for β.

   The $\vee$-elimination rule shows that the following is a derived rule, and hence that $\vee$ induces the *least* upper bound:

   $$\frac{\Gamma, \alpha \vdash_{\mathcal{N}} \gamma \qquad \Gamma, \beta \vdash_{\mathcal{N}} \gamma}{\Gamma, \alpha \vee \beta \vdash_{\mathcal{N}} \gamma}$$

   says that anything above α and β must also be above $\alpha \vee \beta$.

We leave to the reader the trivial check that the operation induced by ∨ is well-defined on equivalence classes, in the sense that

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \leftrightarrow \alpha' \quad \Gamma \vdash_{\mathcal{N}} \beta \leftrightarrow \beta'}{\Gamma \vdash_{\mathcal{N}} (\alpha \vee \beta) \leftrightarrow (\alpha' \vee \beta').}$$

6. $\vdash_{\mathcal{N}}$ *admits no least element*
   There is no formula α such that $\alpha \vdash_{\mathcal{N}} \beta$ for every formula β. Suppose otherwise, and choose a letter p. Then $\alpha \vdash_{\mathcal{N}} p$, and by the Normalization Theorem there is a normal proof of p from α. But since p has no logical symbol, the only possibility is $\alpha = p$, and this should hold for *every* letter p, contradiction.

We can now extend the notion of a Heyting $\sqcap$-algebra.

Definition 17.1.20 (Ogasawara [1939], Birkhoff [1940], McKinsey and Tarski [1946]) *A* **Heyting $\sqcup\sqcap$-algebra** *is a Heyting $\sqcap$-algebra with a l.u.b. operation $\sqcup$.*

As $\sqcap$ produced a lowersemilattice, $\sqcup$ produces now an **uppersemilattice**, i.e. a partially ordered structure in which every pair of elements has a l.u.b. This implies that every non empty finite subset has a l.u.b., but leaves open the degenerate case of the empty set (since the l.u.b. for $\emptyset$ would be the least element). The lack of symmetry between the treatment of ∧ and ∨ expressed by the existence of a greatest element, but not of a least one, is the reason underlying the introduction of ⊥ in the next section.

In the following result we continue the enumeration of equations used in 5.1.7 and 5.1.8.

Proposition 17.1.21 Equational Presentation of Uppersemilattices (Huntington [1904]) *In a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcup \rangle$$

$\sqsubseteq$ is a partial ordering with $=$ as associated equality and $\sqcup$ as associated l.u.b. if and only if

$$x \sqsubseteq y \;\Leftrightarrow\; (x \sqcup y) = y,$$

and the following hold:

9. $x \sqcup x = x$ *(*__idempotency__*)*

10. $x \sqcup (y \sqcup z) = (x \sqcup y) \sqcup z$ *(*__associativity__*)*

11. $x \sqcup y = y \sqcup x$ *(*__commutativity__*).*

The additional condition for 0 being the least element, when it exists, is:

12. $x \sqcup 0 = x$.

Proof. As in 5.1.7. □

We now have equational representations of both g.l.b.'s and l.u.b.'s. But we still have to ensure that the two orderings defined by using $\sqcap$ and $\sqcup$ coincide. The appropriate algebraic notion here is the following.

Definition 17.1.22 *A partially ordered set is a* **lattice** *if it is both a lowersemilattice and an uppersemilattice.*

Proposition 17.1.23 Equational Presentation of Lattices (Huntington [1904]) *In a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, \sqcup \rangle$$

$\sqsubseteq$ *is a partial ordering with* $=$ *as associated equality,* $\sqcap$ *as associated g.l.b. and* $\sqcup$ *as associated l.u.b. if and only if*

$$x \sqsubseteq y \;\Leftrightarrow\; (x \sqcap y) = x \;\Leftrightarrow\; (x \sqcup y) = y,$$

and the following hold, together with 1–3 of 5.1.7 and 9–11 of 17.1.21:

13. $(x \sqcap y) \sqcup y = y$

14. $(x \sqcup y) \sqcap x = x$.

Proof. We only have to deal with the added conditions, since the remaining ones have already been dealt with in 5.1.7 and 17.1.21. They are obviously necessary. Conversely, suppose they hold. We show that

$$(x \sqcap y) = x \;\Leftrightarrow\; (x \sqcup y) = y.$$

If $(x \sqcap y) = x$, then, by 13,

$$x \sqcup y = (x \sqcap y) \sqcup y = y.$$

If $(x \sqcup y) = y$, then, by 14 and commutativity (condition 3 of 5.1.7),

$$x \sqcap y = x \sqcap (x \sqcup y) = x.$$

Thus the two definitions of order coincide. □

Before we proceed further, we notice that Heyting $\sqcup\sqcap$-algebras are lattices of a special kind.

Proposition 17.1.24 *A Heyting* $\sqcup\sqcap$*-algebra is a distributive lattice, i.e. for every* x*,* y *and* z*,*

$$x \sqcap (y \sqcup z) = (x \sqcap y) \sqcup (x \sqcap z)$$

and

$$x \sqcup (y \sqcap z) = (x \sqcup y) \sqcap (x \sqcup z).$$

Proof. Since $\sqcap$ has a right adjoint it must preserve any l.u.b. by 5.3.6, and thus the first property holds. The second one holds symmetrically (by interchanging $\sqcap$ and $\sqcup$ in the proof, since their axioms are symmetrical). □

Exercises 17.1.25 a) *Not every distributive lattice is a Heyting* $\sqcup\sqcap$*-algebra.* (Hint: take any infinite set A and consider the lattice $A^\star$ consisting of A and all its finite subsets, with the usual set theoretical operations. If X is finite, then $X \Rightarrow \emptyset$ does not exist, otherwise it would have to be equal to $A - X$, which is infinite.)

b) *Given a* $\sqcup\sqcap$*-Heyting algebra* $\mathcal{A}$ *and a filter* F *on it, the quotient* $\mathcal{A}_{/F}$ *defined in 5.1.13 is a* $\sqcup\sqcap$*-Heyting algebra.* (Hint: to prove that $\sqcup$ induces an *upper bound*, prove the following *covariance property*:

$$\text{if } a \sqsubseteq b, \text{ then } (c \Rightarrow a) \sqsubseteq (c \Rightarrow b).$$

And to prove that $\sqcup$ induces the *least* upper bound, prove that

$$(x \Rightarrow z) \sqcap (y \Rightarrow z) = (x \sqcup y \ \Rightarrow z),$$

both times using distributivity.)

We now proceed to apply the notion of a Heyting $\sqcup\sqcap$-algebra in the usual way. Definition 5.2.1 of an algebraic consequence refers only to $\sqsubseteq$ and $\sqcap$, and it is unchanged.

Theorem 17.1.26 Algebraic Soundness and Completeness (Jaskowski [1936], Stone [1937], Tarski [1938], McKinsey and Tarski [1948], Rasiowa [1951]) *For any* Γ *and* α*,*

$$\Gamma \vdash_{\mathcal{N}} \alpha \ \Leftrightarrow \ \Gamma \models_a \alpha.$$

Proof. The proof of Soundness is as in 5.2.2, with the additional following cases for disjunction.

If $\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha$ by $\vee$-introduction, then

$$[\![\Gamma]\!]_\rho \sqsubseteq [\![\alpha]\!]_\rho \sqsubseteq ([\![\alpha]\!]_\rho \sqcup [\![\beta]\!]_\rho) = [\![\alpha \vee \beta]\!]_\rho$$

by the induction hypothesis, because $\sqcup$ is a lower bound w.r.t. $\sqsubseteq$, and by definition of $[\![\ \ ]\!]_\rho$. Similarly for β.

If $\Gamma \vdash_{\mathcal{N}} \gamma$ is obtained from $\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta$, $\Gamma, \alpha \vdash_{\mathcal{N}} \gamma$ and $\Gamma, \beta \vdash_{\mathcal{N}} \gamma$ by $\vee$-elimination, then

$$[\![\Gamma]\!]_\rho \sqsubseteq [\![\alpha \vee \beta]\!]_\rho = [\![\alpha]\!]_\rho \sqcup [\![\beta]\!]_\rho$$

by the first induction hypothesis and definition of $[\![\ ]\!]_\rho$, and

$$([\![\Gamma]\!]_\rho \sqcap [\![\alpha]\!]_\rho) \sqsubseteq [\![\gamma]\!]_\rho \qquad \text{and} \qquad ([\![\Gamma]\!]_\rho \sqcap [\![\beta]\!]_\rho) \sqsubseteq [\![\gamma]\!]_\rho$$

by the remaining induction hypotheses. Then

$$\begin{array}{rcl} [\![\Gamma]\!]_\rho & = & [\![\Gamma]\!]_\rho \sqcap ([\![\alpha]\!]_\rho \sqcup [\![\beta]\!]_\rho) \\ & = & ([\![\Gamma]\!]_\rho \sqcap [\![\alpha]\!]_\rho) \sqcup ([\![\Gamma]\!]_\rho \sqcap [\![\beta]\!]_\rho) \\ & \sqsubseteq & [\![\gamma]\!]_\rho \end{array}$$

by the first induction hypothesis, distributivity (17.1.24), and the remaining induction hypotheses.

The proof of Completeness is as in 5.2.3. □

All examples of Heyting $\sqcap$-algebras given in Chapter 4 can be extended to deal with l.u.b.'s in a natural way. For *linear orderings*, the l.u.b. of two elements is their greatest one. For *power sets* and *topological spaces*, it is set-theoretical union. For *lattices*, it is the usual l.u.b. Moreover, *Kripke models* still can be seen as topological Heyting $\sqcup\sqcap$-algebras as in 5.3.10.

Exercise 17.1.27 *Beth models can be seen as topological* $\sqcup\sqcap$*-Heyting algebras.* (Beth [1956]) (Hint: given A $= \langle P, \sqsubseteq, \{A_\sigma\}\rangle$, we cannot consider the order topology on P, because by 5.3.10 it would agree not with Beth but with Kripke forcing on disjunctions. Rather, we consider the topology on the branches f of P generated by the basic open sets $\{f : f \supseteq \sigma\}$, with $\sigma \in P$. If ρ is the environment defined by

$$\rho(p) = \{f : (\exists \sigma \in P)(f \supseteq \sigma \ \wedge \ \sigma \Vdash_{\mathcal{A}} p)\},$$

then

$$[\![\alpha]\!]_\rho = \{f : (\exists \sigma \in P)(f \supseteq \sigma \ \wedge \ \sigma \Vdash_{\mathcal{A}} \alpha)\},$$

by induction on α.)

It is also a routine matter to extend the Stone Representation Theorem 5.4.2 from lowersemilattices to lattices, along the lines of the proof of Kripke Completeness Theorem 17.1.10.

Proposition 17.1.28 Stone Representation Theorem for Heyting $\sqcup\sqcap$-Algebras (Stone [1937]) *Any Heyting* $\sqcup\sqcap$*-algebra is isomorphic to a subalgebra of a topological Heyting* $\sqcup\sqcap$*-algebra.*

Proof. We refer to the proof of 5.4.2. Instead of considering filters, we consider **saturated filters**, i.e. filters F such that, for any x and y in A,

$$x \sqcup y \in F \implies x \in F \vee y \in F,$$

and define a function f from A to the set $\mathcal{F}$ of all saturated filters on A as follows:

$$f(x) = \text{the set of all saturated filters containing } x.$$

Then the first four cases of the proof of 5.4.2 still hold, the fifth requires a supplement of proof, and a new case is needed for $\sqcup$.

5. f *is one-one*
   We only have to prove that *in a distributive lattice, any filter not containing a can be extended to a saturated filter not containing a.* Then, as in 5.4.2, if $x \neq y$, either $x \not\sqsubseteq y$ or $y \not\sqsubseteq x$. If $x \not\sqsubseteq y$, then the upward closure of x is obviously a filter containing x but not y, and by distributivity (17.1.24) there is a saturated filter containing x but not y, i.e. $f(x) \neq f(y)$. Similarly when $y \not\sqsubseteq x$. Thus
   $$x \neq y \implies f(x) \neq f(y).$$

   To prove the claim above, let F be a filter not containing a, and G be a maximal filter w.r.t. $\subseteq$ including F and not containing a. G exists by Zorn's Lemma, since the union of every chain of filters including F and not containing a is still such. As in 17.1.10, we prove that G is saturated.

   Suppose $x \sqcup y \in G$, but $x \notin G$ and $y \notin G$. By maximality, the filter generated by G and x contains a, and thus there is some $b_1 \in G$ such that $b_1 \sqcup x \sqsubseteq a$. Similarly, there is $b_2 \in G$ such that $b_2 \sqcup y \sqsubseteq a$. By letting $b = b_1 \sqcup b_2$,
   $$b \sqcap x \sqsubseteq a \qquad \text{and} \qquad b \sqcap y \sqsubseteq a$$
   and, by distributivity,
   $$b \sqcap (x \sqcup y) = (b \sqcap x) \sqcup (b \sqcap y) \sqsubseteq a.$$
   But both b and $x \sqcup y$ are in G, and thus so are $b \sqcap (x \sqcup y)$ (by closure under $\sqcap$) and a (by upward closure), contradiction.

6. f *preserves* $\sqcup$
   Given x, y and a saturated filter F, if F contains x or y, then it also contains $x \sqcup y$ (by upward closure, since $x \sqsubseteq x \sqcup y$ and $y \sqsubseteq x \sqcup y$), and so $f(x) \cup f(y) \subseteq f(x \sqcup y)$.

   Conversely, if F contains $x \sqcup y$, then it contains one of x and y by saturation, and so $f(x \sqcup y) \subseteq f(x) \cup f(y)$. Thus
   $$f(x \sqcup y) = f(x) \cup f(y). \quad \Box$$

By simply forgetting about topologies, the proof just given shows that *any distributive lattice is isomorphic to a sublattice of a power set, i.e. to a lattice of sets* (Birkhoff [1933]).

Exercises 17.1.29 Prime and maximal filters. A filter in a lattice is **prime** if it contains x or y, whenever it contains $x \sqcup y$. And it is **maximal** if it is nontrivial, and there is no nontrivial filter properly extending it.

a) *In a distributive lattice a maximal filter is prime.* (Hint: suppose F is maximal, $x \sqcup y \in F$ and $x \notin F$. By maximality, the filter generated by $F \cup \{x\}$ is trivial, and y belongs to it. By distributivity, there is $a \in F$ such that $a \sqcup x \sqsubseteq y$, and

$$y = (a \sqcup x) \sqcup y = (a \sqcup y) \sqcap (a \sqcup y).$$

But $x \sqcup y \in F$ by hypothesis, and $a \sqcup y \in F$ because $a \in F$. Then $y \in F$.)

b) *In a non distributive lattice, a maximal filter is not necessarily prime.* (Hint: consider the lattice

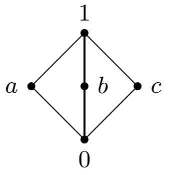

and the filter $\{a, 1\}$. Then $b \sqcup c = 1$ is in it, but neither b nor c are.)

c) *Even in a* $\sqcup\sqcap$*-Heyting algebra, a prime filter is not necessarily maximal.* (Hint: consider the lattice

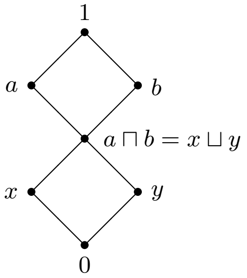

and the filter $\{a, 1\}$.)

Cartesian closed categories with coproducts

The Disjunction Property

The next result exposes a crucial property of $\vee$, typical of intuitionistic logic but not of classical logic.

Theorem 17.1.30 Disjunction Property (Gödel [1933], Gentzen [1935]) *For any* α *and* β,

$$\vdash_{\mathcal{N}} \alpha \vee \beta \ \Rightarrow \vdash_{\mathcal{N}} \alpha \ \textit{or} \ \vdash_{\mathcal{N}} \beta.$$

Proofs. We give a proof for each of the three main systems considered so far.

1. *Natural Deduction*
   Consider a normal proof of $\alpha \vee \beta$ without assumptions, which exists by the Weak Normalization Theorem and the hypothesis $\vdash_{\mathcal{N}} \alpha \vee \beta$. We proceed by induction on the height of the proof. Since there is no assumption, every formula occurring in the proof must be a subformula of $\alpha \vee \beta$. Let us consider the last rule used in the proof.

   If it is an introduction, then it must be a $\vee$-introduction, i.e. one of

   $$\frac{\alpha}{\alpha \vee \beta} \quad \text{or} \quad \frac{\beta}{\alpha \vee \beta.}$$

   Then the proof minus the last step is a proof of either α or β.

   It is now enough to show that the last step of the proof cannot be an elimination. Clearly it cannot be a $\rightarrow$ or $\wedge$ elimination, since they all involve formulas more complicated than their conclusion (while the proof does not, being normal). And it cannot be a $\vee$-elimination either, for suppose it were:

   $$\frac{\begin{matrix} \\ \mathcal{D} \\ \gamma \vee \delta \end{matrix} \qquad \begin{matrix} [\gamma] \\ \mathcal{D}_1 \\ \alpha \vee \beta \end{matrix} \qquad \begin{matrix} [\delta] \\ \mathcal{D}_2 \\ \alpha \vee \beta \end{matrix}}{\alpha \vee \beta.}$$

   On the one hand, any descending path going through $\gamma \vee \delta$ must consist only of eliminations (since the proof is normal, and $\gamma \vee \delta$ is eliminated in the last step), and hence no hypothesis of $\mathcal{D}$ can hence be discharged. On the other hand, the given proof proves $\alpha \vee \beta$ without assumptions, and hence all hypotheses of $\mathcal{D}$ must be discharged. Thus this case cannot happen. □

2. *Sequent System*
   If $\vdash_{\mathcal{S}} \alpha \vee \beta$ then the last step of any proof must be a $\vee$-introduction on the right, and thus the proof minus the last step must be a proof of either $\vdash_{\mathcal{S}} \alpha$ or $\vdash_{\mathcal{S}} \beta$.

3. *Kripke Models*
   Suppose $\not\models_i \alpha$ and $\not\models_i \beta$. Then there is a Kripke model $\mathcal{A}_\alpha$ with least node $0_{\mathcal{A}_\alpha}$ that does not force α. Similarly, there is a Kripke model $\mathcal{A}_\beta$ with least node $0_{\mathcal{A}_\beta}$ that does not force β. We can define a new Kripke model $\mathcal{A}$ with least node $0_{\mathcal{A}}$ by making a disjoint union of $\mathcal{A}_\alpha$ and $\mathcal{A}_\beta$, in such a way that $0_{\mathcal{A}_\alpha}$ and $0_{\mathcal{A}_\beta}$ are the only immediate (and incomparable) extensions of $0_{\mathcal{A}}$. Then $\alpha \vee \beta$ is not forced at $0_{\mathcal{A}}$. If it were, then (by definition of forcing) one of α and β would be forced at $0_{\mathcal{A}}$ and, by monotonicity, at both $0_{\mathcal{A}_\alpha}$ and $0_{\mathcal{A}_\beta}$, contradiction. □

Of course *the result fails, in general, when relativized to sets of premises* Γ. For example, $p \vee q \vdash_{\mathcal{N}} p \vee q$ but $p \vee q \not\vdash_{\mathcal{N}} p$ and $p \vee q \not\vdash_{\mathcal{N}} q$, since the last two sequents are not even classically valid.

We now look for sufficient conditions on Γ that would still allow the result to hold, and the proof of the next result will justify the following definition.

Definition 17.1.31 (Harrop [1960]) *The set of* **Harrop formulas** *is defined inductively as follows:*

1. *a propositional letter is a Harrop formula*
2. *if both α and β are Harrop formulas, so is $\alpha \wedge \beta$*
3. *if β is a Harrop formula then so is $\alpha \to \beta$, for any (not necessarily Harrop) formula α.*

Theorem 17.1.32 Extended Disjunction Property (Harrop [1960]) *If Γ is a set of Harrop formulas, then for every α and β:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta \ \Rightarrow\ \Gamma \vdash_{\mathcal{N}} \alpha \text{ or } \vdash_{\mathcal{N}} \beta.$$

Proof. We consider the simplest among the proofs of the unrelativized result (for $\Gamma = \emptyset$) given above, namely the one using sequents, and proceed by induction on the height of a given proof of $\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta$.

If the last step is an introduction on the right (which, incidentally, must be the case if Γ consists only of propositional letters, i.e. the simplest case of Harrop formulas), then it must be a $\vee$-introduction, and the proof minus the last step must be either a proof of $\Gamma \vdash_{\mathcal{S}} \alpha$ or $\Gamma \vdash_{\mathcal{S}} \beta$.

If the last step is an introduction on the left, there are three possible cases:

1. *conjunction*
   Then $\Gamma = \Gamma_0 \cup \{\gamma \wedge \delta\}$, and e.g.

$$\frac{\Gamma_0, \delta \vdash_{\mathcal{S}} \alpha \vee \beta}{\Gamma_0, \gamma \wedge \delta \vdash_{\mathcal{S}} \alpha \vee \beta.}$$

   We can apply the induction hypothesis because the proof of $\Gamma_0, \delta \vdash_{\mathcal{S}} \alpha \vee \beta$ has smaller height than the proof of $\Gamma_0, \gamma \wedge \delta \vdash_{\mathcal{S}} \alpha \vee \beta$, and if $\gamma \wedge \delta$ is a Harrop formula, then so must be δ. Thus, e.g., $\Gamma_0, \delta \vdash_{\mathcal{S}} \alpha$ and hence

$$\frac{\Gamma_0, \delta \vdash_{\mathcal{S}} \alpha}{\Gamma_0, \gamma \wedge \delta \vdash_{\mathcal{S}} \alpha.}$$

2. *implication*
   Then $\Gamma = \Gamma_0 \cup \{\gamma \to \delta\}$, and

$$\frac{\Gamma_0 \vdash_{\mathcal{S}} \gamma \qquad \Gamma_0, \delta \vdash_{\mathcal{S}} \alpha \vee \beta}{\Gamma_0, \gamma \to \delta \vdash_{\mathcal{S}} \alpha \vee \beta.}$$

   We can apply the induction hypothesis because the proof of $\Gamma_0, \delta \vdash_{\mathcal{S}} \alpha \vee \beta$ has smaller height than the proof of $\Gamma_0, \gamma \to \delta \vdash_{\mathcal{S}} \alpha \vee \beta$, and if $\gamma \to \delta$ is a Harrop formula, then so must be δ. Thus, e.g., $\Gamma_0, \delta \vdash_{\mathcal{S}} \alpha$ and hence

$$\frac{\Gamma_0 \vdash_{\mathcal{S}} \gamma \qquad \Gamma_0, \delta \vdash_{\mathcal{S}} \alpha}{\Gamma_0, \gamma \to \delta \vdash_{\mathcal{S}} \alpha.}$$

3. *disjunction*
   Then $\Gamma = \Gamma_0 \cup \{\gamma \vee \delta\}$, and

$$\frac{\Gamma_0, \gamma \vdash_{\mathcal{S}} \alpha \vee \beta \qquad \Gamma_0, \delta \vdash_{\mathcal{S}} \alpha \vee \beta}{\Gamma_0, \gamma \vee \delta \vdash_{\mathcal{S}} \alpha \vee \beta.}$$

   The proofs of $\Gamma_0, \gamma \vdash_{\mathcal{S}} \alpha \vee \beta$ and $\Gamma_0, \delta \vdash_{\mathcal{S}} \alpha \vee \beta$ both have smaller height than the proof of $\Gamma_0, \gamma \vee \delta \vdash_{\mathcal{S}} \alpha \vee \beta$. But even if we could apply the induction hypothesis, then we would only know that one of α and β is deducible in each case, but not necessarily the same one. For example, we could have $\Gamma_0, \gamma \vdash_{\mathcal{S}} \alpha$ and $\Gamma_0, \delta \vdash_{\mathcal{S}} \beta$, from which we could deduce neither $\Gamma, \gamma \vee \delta \vdash_{\mathcal{S}} \alpha$ nor $\Gamma, \gamma \vee \delta \vdash_{\mathcal{S}} \beta$. Thus the case of disjunction must be excluded.

The three cases justify the definition of Harrop formulas: the first two show the need of inductively considering components of conjunctions and consequents of implications; the last one, as well as the counterexample given before the theorem, show the need of excluding disjunctions. □

17.2 Falsity and Negation

The various connectives of propositional logic have been introduced for different reasons. Implication is the basic object of study, because of its connections with arrow types and hence with functions. Conjunction was needed to reduce finite sets of premises to a single one, which allowed a neat description of the properties of → through an adjointness condition. Disjunction was mainly introduced for symmetry, as a connective dual to conjunction. We now have to dispose of one final lack of symmetry, which is exposed by the algebraic framework.

Syntax

Recall that provable equivalence defines an equivalence relation on formulas, on whose equivalence classes $\vdash_{\mathcal{N}}$, $\wedge$ and $\vee$ respectively induce a partial ordering, as well as the greatest lower bound and least upper bound operations. But while there is a greatest element, corresponding to provable formulas, we argued on p. 383 that there could not be a least element.

This is repaired as follows:

1. the **language** has an added constant $\perp$ (*falsity* or *contradiction*)

2. the definition of **formulas** has an added atomic clause, i.e.

   - $\perp$ is a formula.

As usual, each of the systems we have been considering can be extended to deal with the constant $\perp$, with the intended meaning that $\perp$ is the dual of a provable formula.

Definition 17.2.1 *The relation $\vdash_{\mathcal{N}}$ defined in 1.1.1, 4.1.1 and 17.1.1 is extended to $\perp$ as follows:*

8. **$\perp$-elimination**. *If $\perp$ is deducible from Γ, then any formula α is:*

$$\frac{\Gamma \vdash_{\mathcal{N}} \perp}{\Gamma \vdash_{\mathcal{N}} \alpha.}$$

In terms of rules to extend proofs, this corresponds to the step

$$\frac{\perp}{\alpha.}$$

The intuition is that we should be able to derive a contradiction only when something went wrong. Thus falsity has no introduction rule, unlike all the other connectives introduced so far. Moreover, the elimination rule for falsity tells us that if we derived a contradiction from Γ, then we could derive anything from it.

Since there is no $\perp$-introduction rule, $\perp$ cannot be a maximum itself. But *each instance of the $\perp$-elimination rule can be thought of as a new case of an introduction rule* for α, and thus there is now the possibility of having new maxima (when the proof continues with an elimination of the principal connective of α). For example:

$$\cfrac{\cfrac{\perp}{\alpha \wedge \beta}}{\alpha.}$$

Such maxima are easily eliminated in one step, by going directly from $\bot$ to the conclusion of the elimination rule. For example,

$$\text{from} \quad \cfrac{\cfrac{\bot}{\alpha \wedge \beta}}{\alpha} \quad \text{to} \quad \cfrac{\bot}{\alpha.}$$

The following results continue to hold, with minor adjustements in the proofs.

Proposition 17.2.2 Structure of Normal Proofs (Prawitz [1965]) *For a normal proof of $\mathcal{N}$ the following hold:*

1. **Elimination-Introduction Separation**. *Disregarding repetitions of formulas, any descending path consists of two (possibly empty) parts: a first (upper) one going only through elimination rules, and a second (lower) one going only through introduction rules.*
2. **Subformula Property**. *Any formula occurring in the proof is a subformula of either an undischarged assumption or the conclusion.*

In particular, because of the dual nature of the $\bot$-elimination rule, which can also be seen as an introduction rule, applications of such a rule in a normal proof can occur only at the end of the upper part of any path. Or, equivalently, at the beginning of the lower part.

Theorem 17.2.3 Weak Normalization (Prawitz [1965]) *Every proof can be transformed into a normal proof, by means of an appropriate sequence of maxima eliminations.*

Clearly, the appropriate notion of segment is now that of a sequence of occurrences of a formula α that starts with the conclusion of an introduction of α (either through a usual introduction rule, or through a $\bot$-elimination), and ends with the major premise of an elimination rule.

There is no problem in adding the appropriate axiom to Hilbert systems, to take care of $\bot$. The additions required to prove the following result are trivial.

Theorem 17.2.4 Equivalence of Hilbert Systems and Natural Deduction (Gentzen [1935]) *If $\mathcal{H}$ is any Hilbert system whose theorems include 1–9 of 1.2.3, 4.1.4 and 17.1.4 and, for any α, the following:*

10. $\bot \to \alpha$,

then, for any Γ and β:

$$\Gamma \vdash_{\mathcal{H}} \beta \;\Leftrightarrow\; \Gamma \vdash_{\mathcal{N}} \beta.$$

Finally, the additions to the Sequent System are also trivial.

Definition 17.2.5 (Gentzen [1935]) *The relation $\vdash_{\mathcal{S}}$ defined in 1.3.1, 4.1.5 and 17.1.5 is extended to $\bot$ as follows:*

8. **$\bot$-Rule**. *For every Γ and α,*

$$\frac{\Gamma \vdash_{\mathcal{S}} \bot}{\Gamma \vdash_{\mathcal{S}} \alpha.}$$

Since the new rule for $\bot$ is the same for both $\mathcal{N}$ and $\mathcal{S}$, the two extensions remain equivalent, i.e.

$$\Gamma \vdash_{\mathcal{N}} \beta \;\Leftrightarrow\; \Gamma \vdash_{\mathcal{S}} \beta.$$

Negation ⋆

In intuitionistic logic, negation can be defined using $\bot$.

Definition 17.2.6 *The* **negation** *$\neg\alpha$ of a formula α is defined as $\alpha \to \bot$, i.e. as the assertion that α leads to a contradiction.*

Being a defined symbol, negation does not require rules of its own, but it is sometimes convenient to explicitly state such rules. For example, in the case of natural deduction the following are derived rules, in the usual form of introduction and elimination:

$$\frac{\begin{array}{c}\Gamma, [\alpha]\\ \mathcal{D}\\ \bot\end{array}}{\neg\alpha} \qquad \text{and} \qquad \frac{\alpha \qquad \neg\alpha}{\bot.}$$

On the other hand, we could choose the opposite approach and take *negation as primitive*, with $\bot$ defined as $\alpha \wedge \neg\alpha$ for any α. Rules for $\neg$, independent of the rules for $\bot$, can be introduced as follows, again in the form of introduction and elimination:

$$\frac{\begin{array}{cc}\Gamma, [\alpha] & \Gamma, [\alpha]\\ \mathcal{D}_1 & \mathcal{D}_2\\ \beta & \neg\beta\end{array}}{\neg\alpha} \qquad \text{and} \qquad \frac{\alpha \qquad \neg\alpha}{\beta.}$$

In the usual format for $\mathcal{N}$, these rules would be expressed as:

- **$\neg$-Introduction**. *If both β and $\neg\beta$ are deducible from Γ and α, then $\neg\alpha$ is deducible from Γ:*

$$\frac{\Gamma, \alpha \vdash_{\mathcal{N}} \beta \qquad \Gamma, \alpha \vdash_{\mathcal{N}} \neg\beta}{\Gamma \vdash_{\mathcal{N}} \neg\alpha.}$$

- **¬-Elimination**. *If both α and $\neg\alpha$ are deducible from Γ, then so is any β:*

$$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma \vdash_{\mathcal{N}} \neg\alpha}{\Gamma \vdash_{\mathcal{N}} \beta.}$$

The introduction rule is a restatement of the fact that if α derives a contradiction, then $\neg\alpha$ holds. Similarly, the elimination rule is a restatement of the fact that from a contradiction anything can be derived. This approach is thus not very different from the one above. Moreover, the advantage of having both an introduction and an elimination rule, in the usual spirit of Natural Deduction, is only apparent: the new introduction rule is not completely satisfactory, since the symbol ¬ to be introduced already appears in the premises necessary for its introduction.

Axioms for negation for a Hilbert system simply translate the previous rules:

- $(\alpha \to \beta) \to [(\alpha \to \neg\beta) \to \neg\alpha)$
- $\alpha \to (\neg\alpha \to \beta)$.

Rules for negation are quite natural in the context of Sequent Systems. In this case $\perp$ can be replaced by the empty set, so that the $\perp$-elimination rule becomes a case of a **Thinning Rule on the right**:

$$\frac{\Gamma \vdash_{\mathcal{S}}}{\Gamma \vdash_{\mathcal{S}} \alpha.}$$

We can then formulate the rules for negation quite elegantly and symmetrically, as follows:

- **¬-Introduction on the right**. *If Γ and α are contradictory, then Γ derives $\neg\alpha$:*

$$\frac{\Gamma, \alpha \vdash_{\mathcal{S}}}{\Gamma \vdash_{\mathcal{S}} \neg\alpha.}$$

- **¬-Introduction on the left**. *If Γ derives α, then Γ and $\neg\alpha$ are contradictory:*

$$\frac{\Gamma \vdash_{\mathcal{S}} \alpha}{\Gamma, \neg\alpha \vdash_{\mathcal{S}} .}$$

The advantage of this formulation is that it really introduces ¬ independently of $\perp$. Moreover, $\perp$ *is identified with the empty conclusion*, thus exposing the symmetric roles of empty premises and empty conclusion, which give respectively rise to theorems and contradiction. This approach is particularly neat when used for a presentation of the Classical Propositional Calculus, as we will see in Chapter **??**.

Kripke and Beth models

Kripke and Beth forcing can be extended naturally to the case(s) of $\perp$ (and $\neg$).

Definition 17.2.7 Forcing (Cohen [1963], Kripke [1963]) *For a given possible world $\mathcal{A}$, the relation $\Vdash_{\mathcal{A}}$ defined in 17.1.9 and 17.1.12 is extended to falsity as follows:*

$$\textit{not}\ (\sigma \Vdash_{\mathcal{A}} \perp).$$

In other words, no state forces $\perp$. By the definition of negation, we then have:

$$\begin{aligned} \sigma \Vdash_{\mathcal{A}} \neg\alpha &\Leftrightarrow \sigma \Vdash_{\mathcal{A}} (\alpha \rightarrow \perp) \\ &\Leftrightarrow (\forall \tau \sqsupseteq_{\mathcal{A}} \sigma)(\tau \Vdash_{\mathcal{A}} \alpha \Rightarrow \tau \Vdash_{\mathcal{A}} \perp) \\ &\Leftrightarrow (\forall \tau \sqsupseteq_{\mathcal{A}} \sigma) \text{ not } (\tau \Vdash_{\mathcal{A}} \alpha). \end{aligned}$$

The next result shows that the extension of forcing to $\perp$ captures the intended meaning of contradiction.

Theorem 17.2.8 Kripke Soundness and Completeness (Kripke [1963]) *For any Γ and α:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \Leftrightarrow \Gamma \models_i \alpha.$$

Proof. For the Soundness direction, we supplement the proof of 2.2.5 by the case dealing with $\perp$. If $\Gamma \vdash_{\mathcal{N}} \alpha$ is obtained from $\Gamma \vdash_{\mathcal{N}} \perp$ by $\perp$- *elimination*, then $\Gamma \models_i \perp$ by the induction hypothesis. If σ were any state that forces all formulas of Γ in some world $\mathcal{A}$, then σ would force $\perp$, which is impossible by definition of forcing. Then no σ forces all formulas of Γ, and hence $\Gamma \models_i \alpha$ holds trivially.

For the Completeness direction, we supplement the proofs of 2.2.6 and 17.1.10 (to which we refer) by the case of $\perp$. We have to prove that

$$\Theta \Vdash_{\mathcal{A}} \perp \Leftrightarrow \perp \in \Theta,$$

where Θ is a saturated set of formulas closed under $\vdash_{\mathcal{N}}$. Since the left-hand-side is always false by definition of forcing, we also need the right-hand-side to be always false, and hence we need $\perp \notin \Theta$. This condition can easily be obtained by restricting attention to *nontrivial* or *consistent sets* of formulas, defined precisely by the condition that

$$\perp \notin \Theta.$$

What we have just proved is then that, if $\mathcal{A}$ is the world defined as follows:

$$\mathcal{A} = \langle \mathcal{F}, \subseteq, \{\mathcal{A}_\Theta\}_{\Theta \in \mathcal{F}} \rangle,$$

where:

1. $\mathcal{F}$ is the set of all consistent saturated sets of formulas Θ closed under $\vdash_{\mathcal{N}}$
2. $\subseteq$ is the usual inclusion relation
3. $\mathcal{A}_{\Theta}$ is the set of formulas in Θ consisting only of a propositional letter,

then, for any formula α,

$$\Theta \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \alpha \in \Theta.$$

It remains to finish the proof as usual. We only need to show that *if* $\Gamma \not\vdash_{\mathcal{N}} \alpha$, *then there is a consistent saturated set* Θ *closed under* $\vdash_{\mathcal{N}}$ *such that* $\Theta \supseteq \Gamma$ *and* $\alpha \notin \Theta$ (since then Θ forces every formula in Γ but not α). In 17.1.10 we proved the condition without the consistency part. But $\alpha \notin \Theta$ automatically ensures consistency, and thus that proof suffices. Indeed, if Θ were inconsistent, i.e. $\bot \in \Theta$, then $\Theta \vdash_{\mathcal{N}} \bot$, and $\Theta \vdash_{\mathcal{N}} \alpha$ by the $\bot$-elimination rule. Then $\alpha \in \Theta$ by closure under $\vdash_{\mathcal{N}}$, contradiction. □

Theorem 17.2.9 Beth Soundness (Beth [1956]) *For any* Γ *and* α*:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \;\Leftrightarrow\; \Gamma \models_{ib} \alpha.$$

Proof. For the Soundness direction, we can repeat what has been said above for Kripke forcing.

For the Completeness direction, we supplement the construction of 17.1.14 by the case of $\bot$. There we proved

$$\Gamma_{\sigma} \Vdash_{\mathcal{A}} \alpha \;\Leftrightarrow\; \Gamma_{\sigma} \vdash_{\mathcal{N}} \alpha$$

by induction on α, and we now have to consider the additional case

$$\Gamma_{\sigma} \Vdash_{\mathcal{A}} \bot \;\Leftrightarrow\; \Gamma_{\sigma} \vdash_{\mathcal{N}} \bot.$$

Since the left-hand-side is always false by definition of forcing, we need to have $\Gamma_{\sigma} \not\vdash_{\mathcal{N}} \bot$.

This can be ensured by modifying the construction as follows: for any σ, $\Gamma_{\sigma * \langle 0 \rangle}$ is defined as in 17.1.14 if $\bot$ is not deducible from it, and as Γ_{σ} otherwise. Similarly for $\Gamma_{\sigma * \langle 1 \rangle}$. This takes care of all cases, except for $\Gamma_{\emptyset}$. If Γ itself is consistent, we can let $\Gamma_{\emptyset} = \Gamma$. And if Γ is not consistent, then we know that. On the one hand, $\Gamma \vdash_{\mathcal{N}} \alpha$ holds for every α by $\bot$-elimination. On the other hand, $\Gamma \models_{ib} \alpha$ by the Soundness part proved above. Thus the two sides are still equivalent, and always true. □

The Constructive Model Property can be proved as in 17.1.16, but this time there seems to be no way of avoiding the appeal to decidability of $\vdash_{\mathcal{N}}$: we simply do not want nodes that force $\bot$, and have to know whether they do right away. This causes a problem when one tries to extend the Constructive Model Property to predicate logic, where $\vdash_{\mathcal{N}}$ is not decidable, unless one is willing to relax the definition of Beth models and accept *fallible nodes* that do force $\bot$.

Intuitionistic tableaux

There is no problem in extending the notion of intuitionistic tableau to the case(s) of $\bot$ (and $\neg$).

Definition 17.2.10 *An* **intuitionistic tableau** *is a tree with nodes consisting of signed forcing assertions of the form $T\sigma \Vdash \alpha$ or $F\sigma \Vdash \alpha$, and consistent with the formation rules of 2.3.1, 4.2.4, and 17.1.17, as well as with the following:*

7. *We can extend any branch by adding $F\sigma \Vdash \bot$. Graphically,*

$$\frac{}{\overline{F\sigma \Vdash \bot,}}$$

where the double line and the lack of a top node show that the bottom node can follow any node.

Since the new rule for $\bot$ mirrors the rule for forcing, the two extensions remain equivalent, i.e.

$$\Gamma \vdash_{\mathcal{T}} \alpha \ \Leftrightarrow\ \Gamma \models_i \alpha.$$

From the tableaux rules for $\bot$ and $\to$ (or, directly, from the forcing rule for $\neg$) we derive the following rules for $\neg$:

- *If a node $T\sigma \Vdash \neg\alpha$ is on the tree, then we can extend any branch going through it by adding $F\tau \Vdash \alpha$, where τ is any extension of σ that has already been introduced. Graphically,*

$$\frac{T\sigma \Vdash \neg\alpha}{\overline{F\tau \Vdash \alpha.}}$$

- *If a node $F\sigma \Vdash \neg\alpha$ is on the tree, then we can extend any branch going through it by adding $T\tau \Vdash \alpha$, where τ is a new extension of σ (incomparable with all other extensions of σ already introduced on the same branch). Graphically,*

$$\frac{F\sigma \Vdash \neg\alpha}{\overline{T\tau \Vdash \alpha.}}$$

As usual, the double line shows that the bottom nodes do not have to immediately follow the top one.

Heyting algebras

The addition of $\bot$ requires an algebraic interpretation in terms of a constant.

Definition 17.2.11 Canonical Interpretation. *Given a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, \sqcup, \Rightarrow, 0, 1\rangle$$

and an **environment** ρ *on it, the canonical interpretation* $[\![\ \ ]\!]_\rho$ *defined in 5.1.1 and 17.1.19 is extended to falsity as follows:*

$$[\![\bot]\!]_\rho = 0.$$

The definition of an algebraic model refers only to $\sqsubseteq$ and $\sqcap$, and it is unchanged. We already know that the structure of $\mathcal{N}$ induces a structure of a Heyting $\sqcup\sqcap$-algebra on the equivalence classes of formulas under provable equivalence. We now consider the additional conditions imposed by the rule for $\bot$, by resuming the discussion of pp. 66 and 382, and continuing its enumeration.

7. $\vdash_\mathcal{N}$ *admits a least element*
   The $\bot$-elimination rule
   $$\frac{\Gamma \vdash_\mathcal{N} \bot}{\Gamma \vdash_\mathcal{N} \alpha}$$
   says that what is less than $\bot$ must be less than everything.

This prompts the following extension of the notion of a Heyting $\sqcup\sqcap$-algebra.

Definition 17.2.12 (Ogasawara [1939], Birkhoff [1940], McKinsey and Tarski [1946]) *A* **Heyting algebra** *is a Heyting* $\sqcup\sqcap$*-algebra with a least element 0.*

It is a routine matter to extend the proof of the following result.

Theorem 17.2.13 Algebraic Soundness and Completeness (Jaskowski [1936], Tarski [1937]) *For any* Γ *and* α*,*

$$\Gamma \vdash_\mathcal{N} \alpha \ \Leftrightarrow\ \Gamma \models_a \alpha.$$

Proof. The proof of Soundness is as in 5.2.2, with the additional following case for $\bot$. If $\Gamma \vdash_\mathcal{N} \alpha$ is obtained from $\Gamma \vdash_\mathcal{N} \bot$ by $\bot$-elimination, then

$$[\![\Gamma]\!]_\rho \sqsubseteq [\![\bot]\!]_\rho = 0 \sqsubseteq [\![\alpha]\!]_\rho$$

by the induction hypothesis, definition of $[\![\ \ ]\!]_\rho$, and because 0 is the least element w.r.t. $\sqsubseteq$.

The proof of Completeness is as in 5.2.3. □

All examples of Heyting $\sqcap$-algebras given in Chapter 4 can be extended to deal with the least element in a natural way. For *linear orderings* and *lattices*, its existence must be postulated. For *power sets* and *topological spaces* is the emptyset, which belongs to any topology by definition. Moreover, *Kripke models* still can be seen as topological Heyting algebras.

Exercise 17.2.14 *A Heyting algebra is a linear ordering if and only if, for every pair of elements a and b, $(a \Rightarrow b) \sqcup (b \Rightarrow (b \Rightarrow a)) = 1$.*

It is also a routine matter to extend the Stone Representation Theorem 5.4.2 and 17.1.28, along the lines of the proof of Kripke Completeness Theorem 17.2.8.

Proposition 17.2.15 Stone Representation Theorem for Heyting Algebras (Stone [1937]) *Any Heyting algebra is isomorphic to a subalgebra of a topological Heyting algebra.*

Proof. We refer to the proof of 17.1.28. Instead of considering saturated filters, we consider **nontrivial saturated filters**, i.e. saturated filters F such that $0 \notin F$, and define a function f from A to the set $\mathcal{F}$ of all nontrivial saturated filters on A as follows:

$$f(x) = \text{the set of all nontrivial saturated filters containing } x.$$

Then the six cases of the proofs of 5.4.2 and 17.1.28 still hold. In particular, when proving that in a distributive lattice any filter not containing a can be extended to a saturated filter not containing a, we get nontriviality for free (since if a filter contains 0, then it contains every element by upward closure).

It thus only remains to prove the additional case needed for $\perp$.

7. *f preserves 0*
   Since 0 is in no nontrivial filter,

$$f(0) = \emptyset. \quad \Box$$

Exercise 17.2.16 Pseudocomplements and Ultrafilters. In a Heyting algebra, the **pseudocomplement** of a is the greatest element x such that $a \sqcap x = 0$. And an **ultrafilter** is a nontrivial filter which contains, for any a, either a or its pseudocomplement.

a) *The pseudocomplement of a is $a \Rightarrow 0$.* (Hint: prove that $a \sqcap (a \Rightarrow 0) = 0$, and that if $a \sqcap x = 0$ then $x \sqsubseteq (a \Rightarrow 0)$.)

b) *A non trivial filter is an ultrafilter if and only if it is maximal.* (Hint: if F is a maximal filter and $a \Rightarrow 0 \notin F$, then the filter generated by F and a must be nontrivial, otherwise 0 would belong to it and $a \sqcap x = 0$ for some $x \in F$, in which case $x \sqsubseteq (a \Rightarrow 0)$ and $a \Rightarrow 0 \in F$. By maximality, the filter generated by F and a is equal to F, i.e. $a \in F$.

If F is an ultrafilter and $a \notin F$, then the filter generated by F and a must be trivial because $a \Rightarrow 0 \in F$. Then F is maximal, because no element can be added to it without collapsing it.)

By part b), the consistent and complete sets of formulas in a Lindenbaum algebra are actually ultrafilters.

Bi-cartesian closed categories

The Disjunction Property

There is no problem in extending the notion of a Harrop formula defined in 17.1.31, and using it to extend the Disjunction Property 17.1.32.

Definition 17.2.17 (Harrop [1960]) *The set of* **Harrop formulas** *is defined inductively as follows:*

1. *propositional letters and $\bot$ are Harrop formulas*
2. *if both α and β are Harrop formulas, so is $\alpha \wedge \beta$*
3. *if β is a Harrop formula then so is $\alpha \to \beta$, for any (not necessarily Harrop) formula α.*

Notice that the introduction of $\bot$ as a Harrop formula produces, together with the clause for $\to$ and the definition of $\neg$, the following derived clause:

4. if α is a Harrop formula, so is $\neg\alpha$.

Theorem 17.2.18 Extended Disjunction Property (Harrop [1960]) *If Γ is a set of Harrop formulas, then for every α and β:*

$$\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta \;\Rightarrow\; \Gamma \vdash_{\mathcal{N}} \alpha \textit{ or } \vdash_{\mathcal{N}} \beta.$$

Proof. We only have to supplement the proof of 17.1.32 by the case in which the last step of the proof of $\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta$ is an instance of the $\bot$-rule, i.e.

$$\frac{\Gamma \vdash_{\mathcal{S}} \bot}{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta.}$$

In this case we can obviously derive both $\Gamma \vdash_{\mathcal{S}} \alpha$ and $\Gamma \vdash_{\mathcal{S}} \beta$, again by applying the $\bot$-rule to the premise $\Gamma \vdash_{\mathcal{S}} \bot$. □

In other words, no restriction is needed in the proof because of the presence of $\bot$, and this justifies taking $\bot$ as a Harrop formula.

17.3 A Global Look

We can now take a global look at the full system with all the connectives introduced so far.

Intuitionistic Propositional Calculus

The **language** consists of:

- propositional letters p, q, r, ...
- a constant $\bot$ (*falsity*)
- parentheses '(' and ')'
- the connectives $\rightarrow$ (*implication*), $\wedge$ (*conjunction*), $\vee$ (*disjunction*).

Formulas are defined inductively as follows:

- propositional letters are formulas
- $\bot$ is a formula
- if α and β are formulas, so are $(\alpha \rightarrow \beta)$, $(\alpha \wedge \beta)$, $(\alpha \vee \beta)$.

Natural Deduction

Definition 17.3.1 Natural Deduction (Gentzen [1935]) *The relation $\vdash_{\mathcal{N}}$ has been inductively defined in 1.1.1, 4.1.1, 17.1.1 and 17.2.1 as follows:*

- *assumptions*

  1. $\Gamma, \beta \vdash_{\mathcal{N}} \beta.$

- *implication*

  2. $\dfrac{\Gamma, \alpha \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \rightarrow \beta.}$

  3. $\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma \vdash_{\mathcal{N}} \alpha \rightarrow \beta}{\Gamma \vdash_{\mathcal{N}} \beta.}$

- *conjunction*

  4. $\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \qquad \Gamma \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta.}$

  5. $\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta}{\Gamma \vdash_{\mathcal{N}} \alpha}$ *and* $\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha \wedge \beta}{\Gamma \vdash_{\mathcal{N}} \beta.}$

- *disjunction*

  6. $\dfrac{\Gamma \vdash_{\mathcal{N}} \alpha}{\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta}$ *and* $\dfrac{\Gamma \vdash_{\mathcal{N}} \beta}{\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta.}$

7. $$\frac{\Gamma \vdash_{\mathcal{N}} \alpha \vee \beta \qquad \Gamma, \alpha \vdash_{\mathcal{N}} \gamma \qquad \Gamma, \beta \vdash_{\mathcal{N}} \gamma}{\Gamma \vdash_{\mathcal{N}} \gamma.}$$

- *falsity*

8. $$\frac{\Gamma \vdash_{\mathcal{N}} \bot}{\Gamma \vdash_{\mathcal{N}} \alpha.}$$

Hilbert systems

Recall that, according to definition 1.2.1, $\vdash_{\mathcal{H}}$ is completely determined by the axioms.

Definition 17.3.2 Hilbert System (Herbrand [1928], Tarski [1930], Gentzen [1935]) *The axioms for $\vdash_{\mathcal{H}}$ have been defined in 1.2.3, 4.1.4, 17.1.4 and 17.2.4 as follows, for any α, β, γ and δ:*

- *implication*

  1. $\alpha \to \alpha$

  2. $\gamma \to (\alpha \to \gamma)$

  3. $[(\alpha \to (\gamma \to \delta)] \to [(\alpha \to \gamma) \to (\alpha \to \delta)]$

- *conjunction*

  4. $\alpha \to (\beta \to \alpha \wedge \beta)$

  5. $\alpha \wedge \beta \to \alpha$

  6. $\alpha \wedge \beta \to \beta$

- *disjunction*

  7. $\alpha \to \alpha \vee \beta$

  8. $\beta \to \alpha \vee \beta$

  9. $(\alpha \to \gamma) \to [(\beta \to \gamma) \to (\alpha \vee \beta \to \gamma)]$

- *falsity*

  10. $\bot \to \alpha.$

Sequents

Definition 17.3.3 Sequent System (Gentzen [1935]) *The relation $\vdash_{\mathcal{S}}$ has been inductively defined in 1.3.1, 4.1.5, 17.1.5 and 17.2.5 as follows:*

- *assumptions*

  1. $\Gamma, \beta \vdash_{\mathcal{S}} \beta.$

- *implication*

  2. $\dfrac{\Gamma, \alpha \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \to \beta.}$

  3. $\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \to \beta \vdash_{\mathcal{S}} \gamma.}$

- *conjunction*

  4. $\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha \qquad \Gamma \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \wedge \beta.}$

  5. $\dfrac{\Gamma, \alpha \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \wedge \beta \vdash_{\mathcal{S}} \gamma}$ *and* $\dfrac{\Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \wedge \beta \vdash_{\mathcal{S}} \gamma.}$

- *disjunction*

  6. $\dfrac{\Gamma \vdash_{\mathcal{S}} \alpha}{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta}$ *and* $\dfrac{\Gamma \vdash_{\mathcal{S}} \beta}{\Gamma \vdash_{\mathcal{S}} \alpha \vee \beta.}$

  7. $\dfrac{\Gamma, \alpha \vdash_{\mathcal{S}} \gamma \qquad \Gamma, \beta \vdash_{\mathcal{S}} \gamma}{\Gamma, \alpha \vee \beta \vdash_{\mathcal{S}} \gamma.}$

- *falsity*

  8. $\dfrac{\Gamma \vdash_{\mathcal{S}} \bot}{\Gamma \vdash_{\mathcal{S}} \alpha.}$

Kripke models

Recall that, according to definition 2.2.4, $\models_i$ is completely determined by the notion of Kripke forcing.

Definition 17.3.4 Kripke Forcing (Cohen [1963], Kripke [1963]) *For a given possible world $\mathcal{A}$, the relation of Kripke forcing $\Vdash_{\mathcal{A}}$ has been inductively defined in 2.2.2, 4.2.1, 17.1.9 and 17.2.7 as follows:*

$$\begin{array}{lcl}
\sigma \Vdash_{\mathcal{A}} p & \Leftrightarrow & p \in \mathcal{A}_\sigma \\
\sigma \Vdash_{\mathcal{A}} \alpha \to \beta & \Leftrightarrow & (\forall \tau \sqsupseteq_{\mathcal{A}} \sigma)(\tau \Vdash_{\mathcal{A}} \alpha \;\Rightarrow\; \tau \Vdash_{\mathcal{A}} \beta) \\
\sigma \Vdash_{\mathcal{A}} \alpha \wedge \beta & \Leftrightarrow & \sigma \Vdash_{\mathcal{A}} \alpha \;\; \textit{and} \;\; \sigma \Vdash_{\mathcal{A}} \beta \\
\sigma \Vdash_{\mathcal{A}} \alpha \vee \beta & \Leftrightarrow & \sigma \Vdash_{\mathcal{A}} \alpha \;\; \textit{or} \;\; \sigma \Vdash_{\mathcal{A}} \beta.
\end{array}$$

Moreover,

$$not\ (\sigma \Vdash_{\mathcal{A}} \bot).$$

Beth models

Similarly, $\models_{ib}$ is completely determined by Beth forcing.

Definition 17.3.5 Beth Forcing (Beth [1956]) *For a given possible world $\mathcal{A}$, the relation of Beth forcing $\Vdash_{\mathcal{A}}$ has been inductively defined in 2.2.2, 4.2.1, 17.1.12 and 17.2.7 as follows:*

$$\begin{array}{lcl} \sigma \Vdash_{\mathcal{A}} p & \Leftrightarrow & (\exists B_\sigma)(\forall \tau \in B_\sigma)(p \in \mathcal{A}_\tau) \\ \sigma \Vdash_{\mathcal{A}} \alpha \to \beta & \Leftrightarrow & (\forall \tau \sqsupseteq_{\mathcal{A}} \sigma)(\tau \Vdash_{\mathcal{A}} \alpha \ \Rightarrow \ \tau \Vdash_{\mathcal{A}} \beta) \\ \sigma \Vdash_{\mathcal{A}} \alpha \wedge \beta & \Leftrightarrow & \sigma \Vdash_{\mathcal{A}} \alpha \ \ and \ \ \sigma \Vdash_{\mathcal{A}} \beta \\ \sigma \Vdash_{\mathcal{A}} \alpha \vee \beta & \Leftrightarrow & (\exists B_\sigma)(\forall \tau \in B_\sigma)(\tau \Vdash_{\mathcal{A}} \alpha \ \ or \ \ \tau \Vdash_{\mathcal{A}} \beta), \end{array}$$

where B_σ is a bar for σ. Moreover,

$$not\ (\sigma \Vdash_{\mathcal{A}} \bot).$$

Intuitionistic tableaux

Similarly, $\vdash_{\mathcal{T}}$ is completely determined by the notion of an intuitionistic tableau, which is itself modelled on Kripke forcing.

Definition 17.3.6 Intuitionistic Tableaux (Hughes and Cresswell [1968], Fitting [1983], Nerode [1990]) *An intuitionistic tableau has been defined in 2.3.1, 4.2.4, 17.1.17 and 17.2.10 as a tree with nodes consisting of signed forcing assertions of the form $T\sigma \Vdash \alpha$ or $F\sigma \Vdash \alpha$, and consistent with the following formation rules:*

- *implication*

$$\begin{array}{c} T\sigma \Vdash \alpha \to \beta \\ \hline\hline F\tau \Vdash \alpha \qquad T\tau \Vdash \beta. \end{array} \qquad\qquad \begin{array}{c} F\sigma \Vdash \alpha \to \beta \\ \hline\hline T\tau \Vdash \alpha \\ \hline F\tau \Vdash \beta. \end{array}$$

- *conjunction*

$$\begin{array}{c} T\sigma \Vdash \alpha \wedge \beta \\ \hline\hline T\sigma \Vdash \alpha \\ \hline T\sigma \Vdash \beta. \end{array} \qquad\qquad \begin{array}{c} F\sigma \Vdash \alpha \wedge \beta \\ \hline\hline F\sigma \Vdash \alpha \qquad F\sigma \Vdash \beta. \end{array}$$

- *disjunction*

$$\begin{array}{c} T\sigma \Vdash \alpha \vee \beta \\ \hline\hline T\sigma \Vdash \alpha \qquad T\sigma \Vdash \beta. \end{array} \qquad\qquad \begin{array}{c} F\sigma \Vdash \alpha \vee \beta \\ \hline\hline F\sigma \Vdash \alpha \\ \hline F\sigma \Vdash \beta. \end{array}$$

- *falsity*

$$\frac{\frac{}{}}{F\sigma \Vdash \bot.}$$

The double line shows that the bottom nodes do not have to immediately follow the top one.

Heyting Algebras

Recall that $\models_a$ is completely determined by the notion of a Heyting algebra.

Definition 17.3.7 (Ogasawara [1939], Birkhoff [1940], McKinsey and Tarski [1946]) *A* **Heyting algebra** *has been defined in 5.1.6, 17.1.20 and 17.2.12 as a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \sqcap, \sqcup, \Rightarrow, 0, 1\rangle$$

such that

1. $\sqsubseteq$ *is a partial ordering with* $=$ *as associated equality*
2. $\sqcap$ *is the g.l.b. operation associated with* $\sqsubseteq$
3. $\sqcup$ *is the l.u.b. operation associated with* $\sqsubseteq$*.*
4. $\Rightarrow$ *is the right adjoint of* $\sqcap$ *w.r.t.* $\sqsubseteq$
5. 0 *is the least element of* A *w.r.t.* $\sqsubseteq$
6. 1 *is the greatest element of* A *w.r.t.* $\sqsubseteq$*.*

An alternative characterization of Heyting algebras is obtained from 5.1.7, 5.1.8, 17.1.21 and 17.1.23.

Proposition 17.3.8 Equational Presentation of Heyting Algebras (Huntington [1904], Monteiro [1955], Rasiowa and Sikorski [1963]) *In a structure*

$$\mathcal{A} = \langle A, \sqsubseteq, =, \Rightarrow, \sqcap, \sqcup, 0, 1\rangle$$

$\sqsubseteq$ *is a partial ordering with* $=$ *as associated equality,* 0 *and* 1 *as least and greatest element,* $\sqcap$ *and* $\sqcup$ *as associated l.u.b. and g.l.b., and* $\Rightarrow$ *as the right adjoint of* $\sqcup$ *w.r.t.* $\sqsubseteq$*, if and only if*

$$x \sqsubseteq y \;\Leftrightarrow\; (x \sqcap y) = x,$$

and the following hold:

- *l.u.b. and greatest element*

1. $x \sqcap x = x$
2. $x \sqcap (y \sqcap z) = (x \sqcap y) \sqcap z$
3. $x \sqcap y = y \sqcap x$
4. $x \sqcap 1 = x$

- *adjointness*

5. $(x \Rightarrow x) = 1$
6. $x \Rightarrow (y \sqcap z) = (x \Rightarrow y) \sqcap (x \Rightarrow z)$
7. $x \sqcap (x \Rightarrow y) = x \sqcap y$
8. $y \sqcap (x \Rightarrow y) = y$

- *g.l.b. and least element*

9. $x \sqcup x = x$
10. $x \sqcup (y \sqcup z) = (x \sqcup y) \sqcup z$
11. $x \sqcup y = y \sqcup x$
12. $x \sqcup 0 = x$

- *coherence of l.u.b. and g.l.b.*

13. $(x \sqcap y) \sqcup y = y$
14. $(x \sqcup y) \sqcap x = x.$

Bicartesian Closed Categories

with equational representation, too

Soundness and Completeness Theorems

We can now state the fundamental result summarizing the various Soundness and Completeness Theorems proved throughout the book.

Theorem 17.3.9 The Eightfold Way of Intuitionistic Propositional Calculus. *The following are equivalent, for any* Γ *and* α*:*

1. $\Gamma \vdash_{\mathcal{N}} \alpha$ *(natural deduction)*
2. $\Gamma \vdash_{\mathcal{H}} \alpha$ *(Hilbert system)*
3. $\Gamma \vdash_{\mathcal{S}} \alpha$ *(sequent system)*
4. $\Gamma \vdash_{\mathcal{T}} \alpha$ *(intuitionistic tableaux)*

5. $\Gamma \models_i \alpha$ *(Kripke semantics)*

6. $\Gamma \models_{ib} \alpha$ *(Beth semantics)*

7. $\Gamma \models_a \alpha$ *(Heyting algebras)*

8. $\Gamma \models_c \alpha$ *(bicartesian closed categories).*

Notice how we have four *syntactical* notions of consequence, indicated by the symbol $\vdash$, and four *semantical* ones, indicated by the symbol $\models$. The rule of thumb to distinguish between the two categories is that the first class is existential, requiring *a* proof of a certain kind, while the second class is universal, requiring truth in *all* structures of a given kind. The boundary is obviously somewhat blurred, as implied by the fact that $\vdash_{\mathcal{T}}$ is a syntactical version of a semantical notion. In any case, the illusion of the apparent outer multeplicity of the various incarnations of the syntactical and semantical approaches is dispelled by the teaching of the Eightfold Way, which reveals their real inner unity and exposes their equivalence and interchangeability.

æ

Chapter 18

Heyting Algebras and Topologies

The goal of this chapter is to take a closer look at Heyting algebras. In Sections 1–6 we introduce a linear spectrum of increasingly comprehensive classes, each of which is characterized in a topological way. As a by-product we obtain algebraic characterizations of the most common topological spaces, according to the following table:

Heyting algebras	**topologies**	**typical examples**
finite	finite	
enough strong co-points	closed under intersection	$\mathcal{P}(\omega)$
algebraic	generated by compact opens	2^ω (Cantor space)
continuous	locally quasi-compact	$\mathbb{R}$ (Euclidean space)
enough points	arbitrary	ω^ω (Baire space)
arbitrary		$\mathcal{P}^*(\omega)$

In Section 7 we return to the Algebraic Completeness Theorem from the point of view of special Heyting algebras.

The main themes of the chapter are the notions of a point and a co-point, as well as of a Stone space and a Stone topology. We will barely scratch the surface of the subject, and refer to Gierz, Hofmann, Keimel, Lawson, Mislove and Scott [1980], Johnstone [1982], and Vickers [1989] for detailed treatments.

18.1 Finite Heyting Algebras

We start by considering the finite Heyting algebras, which we can easily be characterized algebraically.

Theorem 18.1.1 Algebraic Characterization of Finite Heyting Algebras. *The finite Heyting algebras are exactly the finite distributive lattices.*

Proof. By 5.3.6, a Heyting algebra is $\sqcap\bigsqcup$-distributive, and hence $\sqcap\sqcup$-distributive.

Conversely, a finite $\sqcap\sqcup$-distributive lattice is $\sqcap\bigsqcup$-distributive, because on a finite lattice the infinitary l.u.b.'s coincide with the finitary ones, and thus it is a Heyting algebra by 5.3.5. □

For the topological characterization of finite Heyting algebras, it is convenient to introduce the notion of a co-point. On the one hand, this is a generalization of the notion of an atom for Boolean algebras (i.e. of an element different from 0 and with no elements between it and 0, see 20.3.4), which does not play any significant role for Heyting algebras.[1] On the other hand, this is a paradigm for further generalizations that will play a fundamental role in this present section (see 18.2.1 and 18.3.1).

Definition 18.1.2 *An element $a \neq 0$ of a lattice is a* **co-point** *if, for every x and y,*

$$a \sqsubseteq x \sqcup y \implies (a \sqsubseteq x) \vee (a \sqsubseteq y).$$

A lattice A has **enough co-points** *if, for every x and y in A,*

$$x \neq y \implies cp(x) \neq cp(y),$$

where

$$cp(x) = \{a : a \textit{ co-point } \wedge\ a \sqsubseteq x\}.$$

In other words, if $x \neq y$, then there is a co-point below one of x and y but not below the other.

The set of co-points of a lattice A is indicated by $\mathbf{Cpt}(\boldsymbol{A})$.

Exercises 18.1.3 a) *In a linear ordering, every element $\neq 0$ is a co-point.*

b) *In a finite lattice, an element $\neq 0$ is a co-point if and only if it has only one immediate predecessor.*

c) *In a distributive lattice, an element $\neq 0$ is a co-point if and only if it is $\sqcup$-irreducible, i.e.*

$$a = x \sqcup y \implies (a = x) \vee (a = y).$$

(Hint: suppose a is $\sqcup$-irreducible, and $a \sqsubseteq x \sqcup y$. Then

$$a = a \sqcap (x \sqcup y) = (a \sqcap x) \sqcup (a \sqcap y)$$

[1] A Heyting algebra can be atomic, in the sense that all elements different from 0 bound an atom, in a trivial way. For example, in a linear ordering there can be at most one atom, and when there is, all elements different from 0 bound the same atom: thus the atoms below the elements do not distinguish them.

by distributivity, so either $a = a \sqcap x$ or $a = a \sqcap y$ by irreducibility, and hence $a \sqsubseteq x$ or $a \sqsubseteq y$.)

d) *In a power set, an element is a co-point if and only if it is a singleton, i.e. of the form* $\{x\}$.

e) *A lattice has enough co-points if and only if every element is the l.u.b. of the co-points below it.* (Hint: given $x \neq 0$, if x is not the l.u.b. of the co-points below it, then there is $y \sqsubset x$ above all such co-points. If A has enough co-points, then there is a co-point below one of x and y but not below the other, contradiction.)

f) *Every finite distributive lattice has enough co-points.* (Hint: by part c), the co-points of a finite distributive lattice are exactly the $\sqcup$-irreducible elements. Every element is either $\sqcup$-irreducible or the l.u.b. of the elements below it and hence, inductively, of the $\sqcup$-irreducible elements below it.)

g) *Every linear ordering has enough co-points.* (Hint: by part a).)

h) *Every power set has enough co-points.* (Hint: by part d), since every set is the union of the singletons contained in it.)

We are not directly interested in complete lattices with enough co-points, since they are not automatically Heyting algebras (by the dual of 18.3.4 below, they are, up to isomorphism, exactly the algebras of *closed* sets of topologies).

Finite topologies

Co-points however are a useful tool in the next proof, which provides a paradigm for the proofs of some crucial later results (18.2.4 and 18.3.4). On its turn, the proof is an extension of an analogous proof for Boolean algebras (20.3.7), which could be usefully read at this point as a warm-up.

Theorem 18.1.4 Topological Characterization of Finite Heyting Algebras. *The finite Heyting algebras are, up to isomorphism, exactly the algebras of open sets of finite topologies, i.e. of topologies with finitely many open sets.*

Proof. To show sufficiency, it is enough to note that the algebra of open sets of a topology is a Heyting algebra (5.3.4).

To show necessity, let A be any finite Heyting algebra, and consider the function cp from A to $\mathcal{P}(\mathrm{Cpt}(A))$ defined as follows:

$$cp(x) = \{a : a \text{ co-point } \wedge\ a \sqsubseteq x\}.$$

Also, as in the proof of 5.4.2, consider the topology generated by $cp(A)$ on $\mathrm{Cpt}(A)$. Then cp is automatically a isomorphism of Heyting algebras, for the following reasons:

- $cp(0) = \emptyset$
  By definition, for any co-point a we have $a \neq 0$, and so $a \not\sqsubseteq 0$.

- $cp(1) = \mathrm{Cpt}(A)$
  By definition, for any element a we have $a \sqsubseteq 1$.

- *if* $x \sqsubseteq y$*, then* $cp(x) \subseteq cp(y)$
  If $x \sqsubseteq y$ and $a \sqsubseteq x$, then $a \sqsubseteq y$, i.e. $cp(x) \subseteq cp(y)$.

- $cp(x \sqcap y) = cp(x) \cap cp(y)$
  For any element a,

  $$a \sqsubseteq (x \sqcap y) \iff (a \sqsubseteq x) \wedge (a \sqsubseteq y)$$

  by definition of $\sqcap$.

- $cp(x \sqcup y) = cp(x) \cup cp(y)$
  For any co-point a,

  $$a \sqsubseteq (x \sqcup y) \iff (a \sqsubseteq x) \vee (a \sqsubseteq y).$$

  The right to left implication holds by definition of $\sqcup$, for any element a. For the left to right implication, let a be a co-point and $a \sqsubseteq x \sqcup y$. Then $a \sqsubseteq x$ or $a \sqsubseteq y$ by definition of co-point.

- $cp(x \Rightarrow y) = (cp(x) \Rightarrow cp(y))$
  As in the proof of 5.4.2, this follows from the fact that right adjointness can be presented equationally in terms of $\sqsubseteq$ and $\sqcap$, which are preserved by cp, and the fact that the topology on the power set of co-points of A is the topology generated by $cp(A)$.

- *if A has enough co-points, then cp is one-one*
  This is just a restatement of the definition of having enough co-points.

  Notice that, by 18.1.3.f, the hypothesis that A has enough co-points is indeed satisfied if A is finite.

- *if A is finite, then cp is onto*
  The proofs given above actually show that cp preserves finite g.l.b.'s and l.u.b.'s. Thus $cp(A)$ is always closed under finite intersections and unions, since

  $$cp(x_1) \cap \cdots \cap cp(x_n) = cp(x_1 \sqcap \cdots \sqcap x_n)$$

  and

  $$cp(x_1) \cup \cdots \cup cp(x_n) = cp(x_1 \sqcup \cdots \sqcup x_n),$$

  and finite g.l.b.'s and l.u.b.'s exist in a Heyting algebra. Since A is finite, so is $cp(A)$, and thus $cp(A)$ is actually closed under arbitrary intersections and unions. Thus the topology generated by $cp(A)$ on the power set of the co-points of A coincides with $cp(A)$, and cp is obviously onto it. □

If A is a Heyting algebra with enough co-points, the set $\mathrm{Cpt}(A)$ of co-points of A and its topology generated by $cp(A)$ are respectively called the **Stone space** of A and the **Stone topology** associated with it.

The Heyting Prime Filter Theorem

The main property of co-points, namely:

- if $a \sqsubseteq x \sqcup y$, then $a \sqsubseteq x$ or $a \sqsubseteq y$,

says that the principal filter generated by a is prime.

If we replace the function

$$cp(x) = \{a : a \text{ co-point } \wedge \; a \sqsubseteq x\}$$

by the function

$$f(x) = \{F : F \text{ is a prime filter containing } x\}$$

in the proof of 18.1.4, we still get a homomorphism of Heyting algebras. Moreover, the proof of the condition

if A has enough co-points, then cp is one-one

shows that

if A has enough prime filters, then f is one-one.

That any Heyting algebra has enough prime filters is the content of the so-called **Heyting Prime Filter Theorem**, saying that if $x \not\sqsubseteq y$ on a Heyting algebra, then there is a prime filter containing x but not y. By verifying it, as we did in 17.1.28,[2] we get the proof of 5.4.2, thus showing that the latter is a generalization of the one-one embedding part of 18.1.4 to arbitrary Heyting algebras (not necessarily having enough co-points). More precisely, the role of co-points is taken by prime filters, and the condition that there are enough co-points becomes the condition that there are enough prime filters, which is always satisfied by the Heyting Prime Filter Theorem.

Thus we extend the notions of **Stone space** and **Stone topology** to arbitrary Heyting algebras (not necessarily with enough co-points), by considering the set $\mathcal{F}_A^p$ of prime filters of A and its topology generated by $f(A)$.

[2] The usual proof (17.1.28) of the Heyting Prime Filter Theorem deduces it from the **Heyting Maximal Filter Theorem**, saying that if $x \not\sqsubseteq y$ on a Heyting algebra, then there is a maximal filter containing x but not y.

The strength of the Heyting Prime or Maximal Filter Theorems is discussed, in a dual context, in note 6 on p. 422.

18.2 Complete Heyting Algebras with Enough Strong Co-points ⋆

The proof of 18.1.4 shows that the collection $cp(A)$ of subsets of $\mathrm{Cpt}(A)$ is closed under finite unions and intersections. All $cp(A)$ lacks to be a topology, is closure under arbitrary unions. When A is finite this closure is automatic, since there are only finitely many subsets of $\mathrm{Cpt}(A)$ anyway. When A is infinite things are more complicated, but $cp(A)$ would still be closed under arbitrary unions, by the same proof, if arbitrary $\bigsqcup$ existed and were preserved by cp.

The first condition is easy to achieve, by requiring A to be complete. For the second condition, we look at the proof that cp preserves finite $\sqcup$, and notice that it used the notion of co-point, i.e. $\sqcup$-irreducibility. To achieve the second condition we thus consider an infinitary version of the notion of co-point, already used in 6.3.28 under the name of strong finiteness.

Definition 18.2.1 *An element $a \neq 0$ of a lattice is a* **strong co-point** *if, for every subset X,*

$$a \sqsubseteq \bigsqcup X \implies (\exists x \in X)(a \sqsubseteq x).$$

A lattice A has **enough strong co-points** *if, for every x and y in A,*

$$x \neq y \implies scp(x) \neq scp(y),$$

where

$$scp(x) = \{a : a \text{ strong co-point } \wedge\ a \sqsubseteq x\}.$$

In other words, if $x \neq y$, then there is a strong co-point below one of x and y but not below the other.

The set of strong co-points of a lattice A is indicated by $\mathbf{Scpt}(\boldsymbol{A})$.

Exercises 18.2.2 The next exercises are similar to the ones in 18.1.3, to which one can turn for hints.

a) *In a linear ordering, an element $\neq 0$ is a strong co-point if and only if it has an immediate predecessor.* (Hint: if a has an immediate predecessor b, and X contains only elements $\sqsubset a$, then $\bigsqcup X \sqsubseteq b$, i.e. $\bigsqcup X \sqsubset a$. Conversely, if a has no immediate predecessor, let X be the set of elements $\sqsubset a$: then $\bigsqcup X = a$, but X has no element $\sqsupseteq a$.)

b) *In a finite lattice, an element $\neq 0$ is a strong co-point if and only if it has only one immediate predecessor.*

c) *In a $\sqcap\bigsqcup$-distributive lattice, an element $\neq 0$ is a strong co-point if and only if it is $\bigsqcup$-irreducible, i.e.*

$$a = \bigsqcup X \implies (\exists x \in X)(a = x).$$

d) *In a power set, an element is a strong co-point if and only if it is a singleton, i.e. of the form $\{x\}$.*

e) *A lattice has enough strong co-points if and only if every element is the l.u.b. of the strong co-points below it.*

f) *Every finite distributive lattice has enough strong co-points.*

g) *In a dense linear ordering there are no strong co-points.*

h) *Every power set has enough strong co-points.*

The next result relates the two classes of Heyting algebras introduced in the present and the previous sections, and it is just a restatement of some of the previous exercises.

Proposition 18.2.3 *Every finite Heyting algebra has enough strong co-points, but there are complete Heyting algebras with enough strong co-points that are not finite.*

Proof. To prove the first part, we want to show that every non-zero element of a finite Heyting algebra is the l.u.b. of the strong co-points below it. In a finite lattice $\bigsqcup$ and $\sqcup$ coincide, and hence so do strong co-points and co-points. If in addition the lattice is distributive, then they both coincide with the $\sqcup$-irreducible elements. And in a finite lattice every non-zero element is obviously the l.u.b. of the $\sqcup$-irreducible elements below it.

To prove the second part, we want to exhibit an infinite, complete Heyting algebra with enough strong co-points. The *algebra of all sets of natural numbers* $\mathcal{P}(\omega)$ provides an example, since every singleton is a strong co-point, and every set is the union of the singletons contained in it. □

Topologies closed under arbitrary intersections

The next result provides a topological characterization of complete Heyting algebras with enough strong co-points.

Theorem 18.2.4 Topological Characterization of Complete Heyting Algebras with Enough Strong Co-points (Büchi [1952], Raney [1952]) *The complete Heyting algebras with enough strong co-points are, up to isomorphism, exactly the algebras of open sets of topologies closed under arbitrary intersections.*

Proof. To show sufficiency, first note that the open sets of a topology $\Omega(X)$ on a set X are always closed under arbitrary unions. If they are also closed under arbitrary intersections, then they form a complete Heyting algebra w.r.t. the set-theoretic operations $\bigcup$ and $\bigcap$. It remains to prove the following two facts:

- *a topology closed under arbitrary intersections always has strong co-points*
  For any element $x \in X$, consider

$$(\{x\})_\circ = \text{smallest open set containing } x,$$

which exists because the open sets are closed under arbitrary intersections (it is enough to take the intersection of all open sets containing x). Suppose $(\{x\})_\circ \subseteq \bigcup_{i\in I} A_i$, with A_i open. Since $x \in (\{x\})_\circ$, at least one of the A_i's must contain x. Since A_i is an open set containing x, then by definition $(\{x\})_\circ \subseteq A_i$, because $(\{x\})_\circ$ is the smallest open set containing x.

- *a topology closed under arbitrary intersections always has enough strong co-points*
  Suppose A and B are distinct open sets. There must be an element x in one of them but not in the other, e.g. $x \in B - A$. Then $(\{x\})_\circ \not\subseteq A$ because $x \notin A$, and $(\{x\})_\circ \subseteq B$ because $x \in B$ and B is open. So $(\{x\})_\circ$ is a strong co-point below one of A and B but not below the other.

To show necessity, let A be any complete Heyting algebra with enough strong co-points, and consider the function *scp* from A to $\mathcal{P}(\mathrm{Scpt}(A))$ defined as follows:

$$\mathit{scp}(x) = \{a : a \text{ strong co-point } \wedge\ a \sqsubseteq x\}.$$

Consider also the topology generated by $\mathit{scp}(A)$ on $\mathrm{Scpt}(A)$. As in the proof of 18.1.4, *scp* is automatically a homomorphism of Heyting algebras. Moreover:

- *if A has enough strong co-points, then scp is one-one*
  This is just a restatement of the definition of having enough strong co-points.

- *if A is complete, then scp is onto*
  The proof of 18.1.4 shows that *scp* preserves arbitrary g.l.b.'s and finite l.u.b.'s. Since we are using the notion of strong co-point instead of the notion of co-point, *scp* also preserves arbitrary l.u.b.'s. If A is complete, then arbitrary g.l.b.'s and l.u.b.'s exist, and thus $\mathit{scp}(A)$ is closed under arbitrary unions and intersections. Thus the topology generated by $\mathit{scp}(A)$ on the power set of the strong co-points of A is $\mathit{scp}(A)$ itself, it is closed under arbitrary intersections, and *scp* is obviously onto it. □

Notice that the proof shows in particular that *a complete lattice with enough strong co-points is a Heyting algebra*, since no use of $\Rightarrow$ is made in the proof of the isomorphism.

18.3 Complete Heyting Algebras with Enough Points

The next notion is dual to 18.1.2.

Definition 18.3.1 *An element $a \neq 1$ of a lattice is a* **point** *if, for every x and y,*

$$a \sqsupseteq (x \sqcap y) \implies (a \sqsupseteq x) \vee (a \sqsupseteq y).$$

A lattice A has **enough points** *if, for every x and y in A,*

$$x \neq y \implies p(x) \neq p(y),$$

where

$$p(x) = \{a : a \text{ point } \wedge\ a \not\sqsupseteq x\}.$$

In other words, if $x \neq y$, then there is a point above one of x and y but not above the other.

The set of points of a lattice A is indicated by $\mathbf{Pt}(\boldsymbol{A})$.

Exercises 18.3.2 The next exercises are dual to the ones in 18.1.3, to which one can turn for hints.

a) *In a linear ordering, every element $\neq 1$ is a point.*

b) *In a finite lattice, an element $\neq 1$ is a point if and only if it has only one immediate successor.*

c) *In a distributive lattice, an element $\neq 1$ is a point if and only if it is $\sqcap$-irreducible, i.e.*

$$a = (x \sqcap y) \implies (a = x) \vee (a = y).$$

d) *In a power set, an element is a point if and only if it is the complement of a singleton, i.e. of the form $\overline{\{x\}}$.*

e) *A lattice has enough points if and only if every element is the g.l.b. of the points above it.*

f) *Every finite distributive lattice has enough points.*

g) *Every linear ordering has enough points.*

h) *Every power set has enough points.*

The next result relates the two classes of Heyting algebras introduced in the present and the previous sections.

Proposition 18.3.3 Papert [1959]) *Every complete Heyting algebra with enough strong co-points has enough points, but there are complete Heyting algebras with enough points that have no (strong) co-point.*

Proof. To prove the first part, let $x \not\sqsubseteq y$. Since the strong co-points below x have l.u.b. x, there is a strong co-point a such that $a \sqsubseteq x$ and $a \not\sqsubseteq y$. We consider the set

$$F = \{z : a \sqsubseteq z\},$$

and notice the following properties:

- F is a filter. In particular, F is closed under upward and under $\sqcap$.

- F contains only elements above a. In particular $y \notin F$.

- No element of $\overline{F}$ is above x. This follows from the fact that F is upward closed and $x \in F$.

- Every maximal element in $\overline{F}$ is $\sqcap$-irreducible, and hence a point (by 18.3.2.c). This follows from the fact that F is closed under $\sqcap$.

- If X is contained in $\overline{F}$, then $\bigsqcup X$ is also in $\overline{F}$. Otherwise $a \sqsubseteq \bigsqcup X$, and since a is a strong co-point there is $x \in X$ such that $a \sqsubseteq x$, i.e. $X \cap F \neq \emptyset$.

Let now C be any maximal chain in $\overline{F}$ containing y, which exists by Zorn's Lemma. Then $\bigsqcup C$ is also in $\overline{F}$. Since C is maximal, $\bigsqcup C$ is a maximal element of $\overline{F}$, and hence a point above y (by the choice of C), but not above x (because it is in $\overline{F}$).

To prove the second part, we want to exhibit a complete Heyting algebra with enough points and no co-point. The *algebra of open sets of the euclidean space* $I\!\!R$ provides an example. First, it has enough points because so does every topology (by the first part of the proof of 18.3.4 below). Second, it has no co-point because every open interval is the union of two different open intervals. □

Arbitrary topologies

By the Stone Representation Theorem (5.4.2, 17.1.28 and 17.2.15), every Heyting algebra is isomorphic to a *sub*algebra of a topological algebra. This leaves open the question, answered in the next result, of which Heyting algebras are actually isomorphic to the *full* topological algebra.

From a complementary perspective, every topological algebra is a Heyting algebra. This leaves open the question, also answered in the next result, of which algebraic properties of Heyting algebras are characteristic of the topological algebras.

The proof of the next result is dual to those of 18.1.4 and 18.2.4.

Theorem 18.3.4 Algebraic Characterization of Topologies (Papert [1959]) *The complete Heyting algebras with enough points are, up to isomorphism, exactly the algebras of open sets of topologies.*

Proof. To show sufficiency, first note that the open sets of a topology $\Omega(X)$ on a set X are closed under arbitrary unions, and hence they form a complete Heyting algebra.[3] It remains to prove the following two facts:

[3] Notice that, while the open sets of a topology are closed under arbitrary unions and finite intersections, and hence infinitary l.u.b.'s and finitary g.l.b.'s are the usual set-theoretical unions and intersections, they are not in general closed under arbitrary intersections, and thus an infinitary g.l.b. is only the largest open set contained in the set-theoretical intersection, i.e. the latter's interior.

- *a topology always has points*
  For any element $x \in X$, consider
  $$\left(\overline{\{x\}}\right)^\circ = \text{interior of } \overline{\{x\}} = \text{greatest open set not containing } x,$$
  which exists because the open sets are closed under arbitrary unions (it is enough to take the union of all open sets not containing x). Suppose $\left(\overline{\{x\}}\right)^\circ \supseteq A \cap B$, with A and B open. At least one of A and B must not contain x, otherwise x would be in the intersection too. Suppose e.g. $x \notin A$. Then $A \subseteq \left(\overline{\{x\}}\right)^\circ$ by definition, since $\left(\overline{\{x\}}\right)^\circ$ is the greatest open set not containing x.

- *a topology always has enough points*
  Suppose A and B are distinct open sets. There must be an element x in one of them but not in the other, e.g. $x \in B - A$. Then $\left(\overline{\{x\}}\right)^\circ \supseteq A$ because $x \notin A$ and A is open, and $\left(\overline{\{x\}}\right)^\circ \not\supseteq B$ because $x \in B$. So $\left(\overline{\{x\}}\right)^\circ$ is a point above one of A and B but not above the other.

To show necessity, let A be any complete Heyting algebra with enough points, and consider the function p from A to $\mathcal{P}(\mathrm{Pt}(A))$ defined as follows:

$$p(x) = \{a : a \text{ point } \wedge \ a \not\sqsupseteq x\}.$$

Consider also the topology generated by $p(A)$ on $\mathrm{Pt}(A)$. As in the proof of 5.4.2, p is automatically a isomorphism of Heyting algebras, for the following reasons:

- $p(0) = \emptyset$
  By definition, for any element a we have $a \sqsupseteq 0$.

- $p(1) = \mathrm{Pt}(A)$
  By definition, for any point a we have $a \neq 1$, and so $a \not\sqsupseteq 1$.

- *if* $x \sqsubseteq y$, *then* $p(x) \subseteq p(y)$
  If $x \sqsubseteq y$ and $a \not\sqsupseteq x$, then $a \not\sqsupseteq y$, i.e. $p(x) \subseteq p(y)$.

- $p(x \sqcap y) = p(x) \cap p(y)$
  For any point a,
  $$a \sqsupseteq (x \sqcap y) \iff (a \sqsupseteq x) \vee (a \sqsupseteq y).$$
  The right to left implication holds by definition of $\sqcap$, for any element a. For the left to right implication, let a be a point and $a \sqsupseteq x \sqcap y$. Then $a \sqsupseteq x$ or $a \sqsupseteq y$ by definition of point.

By taking negations,

$$a \not\sqsupseteq (x \sqcap y) \iff (a \not\sqsupseteq x) \wedge (a \not\sqsupseteq y).$$

- $p(x \sqcup y) = p(x) \cup p(y)$
  For any element a,

  $$a \sqsupseteq (x \sqcup y) \iff (a \sqsupseteq x) \wedge (a \sqsupseteq y)$$

  by definition of $\sqcup$. By taking negations,

  $$a \not\sqsupseteq (x \sqcup y) \iff (a \not\sqsupseteq x) \vee (a \not\sqsupseteq y).$$

- $p(x \Rightarrow y) = (p(x) \Rightarrow p(y))$
  As in the proof of 5.4.2, this follows from the fact that right adjointness can be presented equationally in terms of $\sqsubseteq$ and $\sqcap$, which are preserved by p, and the fact that the topology on the power set of points of A is the topology generated by $p(A)$.

- *if A has enough points, then p is one-one*
  This is just a restatement of the definition of having enough points.

- *if A is complete, then p is onto*
  The proofs given above actually show that p preserves finite g.l.b.'s and arbitrary l.u.b.'s. Thus $p(A)$ is always closed under finite intersections, since

  $$p(x_1) \cap \cdots \cap p(x_n) = p(x_1 \sqcap \cdots \sqcap x_n),$$

  and finite g.l.b.'s exist in a Heyting algebra. Similarly, $p(A)$ is closed under arbitrary unions when arbitrary l.u.b.'s exist on A, i.e. when A is a complete Heyting algebra. Thus, if A is complete, the topology generated by $p(A)$ on the power set of the points of A is $p(A)$ itself, and p is obviously onto it. □

Notice that the proof shows in particular that *a complete lattice with enough points is a Heyting algebra*, since no use of $\Rightarrow$ is made in the proof of the isomorphism.

If A is a Heyting algebra with enough points, the set $\mathrm{Pt}(A)$ of points of A and its topology generated by $p(A)$ are respectively called the **dual Stone space** of A and the **dual Stone topology**[4] associated with it.

Exercises 18.3.5 Sober spaces. (Stone [1937], Grothendieck and Dieudonné [1960]) A topological space is called **sober** if:

1. the open sets $\left(\overline{\{x\}}\right)^{\circ}$ are the only points

[4] Some authors call the dual Stone topology the **hull-kernel topology**.

2. if $x \neq y$, then $\left(\overline{\{x\}}\right)^\circ \neq \left(\overline{\{y\}}\right)^\circ$.

A topological space is called $\mathbf{T_2}$ or **Hausdorff** if every two distinct points can be separated by disjoint open sets, i.e. if $x \neq y$, then there are disjoint open sets A and B such that $x \in A$ and $y \in B$.

a) *The two conditions in the definition of a sober space are independent.* (Hint: if X has at least two elements, and $\emptyset$ and X are the only open sets, then 1 holds but 2 fails. If X is infinite and its cofinite subsets are the only nonempty open sets, then 2 holds but 1 fails.)

b) *A sober space is T_0, but there are T_0 not sober spaces.* (Hint: 2 is equivalent to the T_0 axiom. For a T_0 not sober space, see the second example in part a).)

c) *A T_2 space is sober, but there are sober not T_2 spaces.* (Hint: in a T_2 space the points are exactly the complements of singletons, i.e. the ones of the form $\overline{\{x\}}$, because if A is an open set not containing two elements $y \neq z$, then there are disjoint open sets B and C such that $y \in B$ and $z \in C$. Then $A = (A \cup B) \cap (A \cup C)$, and A is not a point.

For a sober not T_2 space, consider the order topology of a finite linear ordering.)

d) *A topology is homeomorphic to the dual Stone topology of a complete Heyting algebra with enough points if and only if it is sober*. (Hint: consider a complete Heyting algebra A with enough points, and its dual Stone topology $\Omega(\mathrm{Pt}(A))$ on the set of points $\mathrm{Pt}(A)$, defined in the proof of 18.3.4. Notice that if X is an open set, then $X = p(a)$ for one and only one a, because p is onto and one-one. Notice that:

- An open set X is a point of the topology if and only if $X = p(a)$ for some point a of A.
- If a is a point of A, then $p(a) = \left(\overline{\{a\}}\right)^\circ$.
  Indeed, if a is a point and $p(x)$ does not contain a, then $a \sqsupseteq x$ by definition of p. So $p(a) \supseteq p(x)$, because p preserves $\sqsupseteq$, and $p(a)$ is the greatest open set not containing a.

From this 1 follows immediately, and 2 follows because p is one-one, i.e. $\Omega(\mathrm{Pt}(A))$ is sober.

Conversely, let X be a space with a sober topology $\Omega(X)$. Then the map

$$x \longmapsto \left(\overline{\{x\}}\right)^\circ \qquad \text{from } X \text{ to } \mathrm{Pt}(\Omega(X))$$

is onto and one-one by 1 and 2, i.e. an isomorphism. And it is automatically a homeomorphism because, being $\Omega(X)$ a complete Heyting algebra with enough points, the map

$$B \longmapsto \{\left(\overline{\{x\}}\right)^\circ : \left(\overline{\{x\}}\right)^\circ \not\supseteq B\} \qquad \text{from } \Omega(X) \text{ to } \Omega(\mathrm{Pt}(\Omega(X)))$$

is an isomorphism.)

The Heyting Prime Ideal Theorem

The reason for the word 'dual' used for the space and topology introduced in the proof of 18.3.4 comes from the fact that we have there switched from the filters used in the previous versions of the Stone Representation Theorem (5.4.2, 17.1.28 and 17.2.15), to ideals.[5] Indeed, the main property of points, namely:

- if $a \sqsupseteq x \sqcap y$, then $a \sqsupseteq x$ or $a \sqsupseteq y$,

says that the principal ideal generated by a is prime.

If we replace the function

$$p(x) = \{a : a \text{ point } \wedge \ a \not\sqsupseteq x\}$$

by the function

$$i(x) = \{I : I \text{ is a prime ideal not containing } x\}$$

in the proof of 18.3.4, we still get a homomorphism of Heyting algebras. Moreover, the proof of the condition

$$\text{if } A \text{ has enough points, then } p \text{ is one-one}$$

shows that

$$\text{if } A \text{ has enough prime ideals, then } i \text{ is one-one.}$$

That any Heyting Algebra has enough prime ideals is the content of the so-called **Heyting Prime Ideal Theorem** saying that if $x \not\sqsubseteq y$ on a Heyting algebra, then there is a prime ideal containing y but not x.[6] By verifying it, with a proof dual to that of 17.1.28, we get a dual proof of 5.4.2.

Thus we extend the notions of **dual Stone space** and **dual Stone topology** to arbitrary Heyting algebras (not necessarily with enough points), by considering the set $\mathcal{I}_A^p$ of prime ideals of A and its topology generated by $i(A)$.

[5] This also accounts for the slight backwardness of the definition of p, since a subset of a lattice is a prime filter if and only if its complement is a prime ideal (see also note 7 on p. 426), and thus prime filters containing x correspond to prime ideals *not* containing x.

Indeed, if we defined

$$p(x) = \{a : a \text{ point } \wedge \ a \sqsupseteq x\},$$

then we would have

$$p(x \sqcap y) = p(x) \cup p(y) \qquad \text{and} \qquad p(x \sqcup y) = p(x) \cap p(y),$$

and thus p would preserve neither of $\sqcap$ and $\sqcup$.

[6] The usual proof of the Heyting Prime Ideal Theorem deduces it from the **Heyting Maximal Ideal Theorem** saying that if $x \not\sqsubseteq y$ on a Heyting algebra, then there is a maximal ideal containing y but not x (see 17.1.28).

The Heyting Prime Ideal Theorem is not provable in ZF (Tarski [1954], Łoš and Ryll-Nardzewski [1954]), it implies a weak form of the Axiom of Choice (namely, the existence of a choice function for any family of nonempty *finite* sets), but it does not imply the full Axiom of Choice (Halpern [1964]).

The Heyting Maximal Ideal Theorem is equivalent to the Axiom of Choice (Scott [1954], Mrowka [1956], and Klimovsky [1958]), and hence it is strictly stronger than the Heyting Prime Ideal Theorem.

While for the sake of the Stone Representation Theorem both filters and ideals eventually produce the same result, we should notice that *filters are more natural from the point of view of logic*, since they correspond to sets of formulas closed under $\vdash_{\mathcal{N}}$ and $\wedge$ (equivalently, by 5.1.11, closed under $\rightarrow$), while *ideals are more natural from the point of view of topology*, since a topology may have no co-point at all (see the proof of 18.3.3), while it always has enough points (see 18.3.4).

Exercises 18.3.6 Points in Heyting algebras of ideals. Given a lattice A, let $\mathcal{I}_A$ be the set of its ideals.

a) *If A is a distributive lattice, then $\mathcal{I}_A$ is a complete Heyting algebra with enough points*. (Hint: the g.l.b. of a finite family of ideals is their set-theoretical intersection $\cap$, and the l.u.b. $\bigsqcup$ of an arbitrary family is the smallest ideal containing the set-theoretical union. Thus $\mathcal{I}_A$ is a complete lattice. $\mathcal{I}_A$ is a Heyting algebra by 5.3.5, since the $\sqcap\bigsqcup$-distributive law holds: if $a \in I \cap (\bigsqcup_{x \in X} J_x)$, then $a \in I$ and $a \in \bigsqcup_{x \in X} J_x$, i.e. $a \sqsubseteq a_1 \sqcup \cdots \sqcup a_n$ for some $a_i \in J_{x_i}$; thus $a \in \bigsqcup_{x \in X}(I \cap J_x)$.)

We now prove that *the points of $\mathcal{I}_A$ are exactly the prime ideals*. Suppose I is a point of $\mathcal{I}_A$, and it contains $a \sqcap b$: then it contains the principal ideal generated by $a \sqcap b$, which is the intersection of the principal ideals generated by a and b; since I is a point, it must contain one of these principal ideals, and hence one of a and b. Conversely, if I is a prime ideal, suppose $I \supseteq I_1 \cap I_2$, but $I \not\supseteq I_1$, i.e. $x \in I_1 - I$ for some x: if $y \in I_2$, then $x \sqcap y$ is in I, and $y \in I$ because I is prime and $x \notin I$; thus $I_2 \subseteq I$.

That $\mathcal{I}_A$ has enough points now follows from the fact that it has enough prime ideals. More precisely, if $I \not\subseteq J$, then there is an $a \in I - J$, and there is a prime ideal containing I but not a, and hence not J. Notice that the existence of a maximal filter containing I but not a is immediate by Zorn's Lemma, and distributivity is needed to deduce that a maximal ideal is prime, see 17.1.28.)

b) *If A is a Heyting algebra, then the dual Stone topologies of A (as an arbitrary Heyting algebra) and $\mathcal{I}_A$ (as a Heyting algebra with enough points) coincide*. (Hint: the dual Stone topology of A is generated by

$$i_1(x) = \{I : I \text{ is a prime ideal} \wedge x \notin I\},$$

while the dual Stone topology of $\mathcal{I}_A$ is generated by

$$i_2(J) = \{I : I \text{ is a prime ideal} \wedge I \not\supseteq J\}.$$

But $I \not\supseteq J$ means that $x \in J - I$ for some x, and hence $i(J) = \bigcup_{x \in J} i(x)$, i.e. a basic open set of the dual Stone topology of $\mathcal{I}_A$ is an open set of the dual Stone topology of A.)

18.4 Arbitrary Heyting Algebras

In the previous section we focused on topologies, and characterized them from the point of view of Heyting algebras. In the present section we focus on arbitrary Heyting algebras, and characterize them from the point of view of topology.

The next result shows that there is indeed a need to look for classes of Heyting algebras more general than those considered so far.

Proposition 18.4.1 *There are complete Heyting algebras without points, as well as Heyting algebras that are not complete.*

Proof. For an example of a complete Heyting algebra without points, consider the *sets of natural numbers modulo cofinite sets* $\mathcal{P}^*(\omega)$, i.e. the quotient Heyting algebra obtained from $\mathcal{P}(\omega)$ w.r.t. the filter of cofinite sets. There is no point because each equivalence class except the greatest one contains coinfinite sets. And if A is coinfinite, then it is the intersection of two coinfinite sets A_1 and A_2 differing coinfinitely from A (e.g. $A_i = A \cup B_i$, where B_1 and B_2 are infinite disjoint sets such that $B_1 \cup B_2 = \overline{A}$).

For an example of a Heyting algebra that is not complete, consider the *set of rational numbers between 0 and 1*, as a linearly ordered set. It is a Heyting algebra by 5.3.1, but is obviously not complete. □

Lindenbaum algebras

In 5.2.2, 5.2.3, 17.1.26 and **??** we have proved the following result.

Theorem 18.4.2 Algebraic Soundness and Completeness (Jaskowski [1936], Tarski [1937]) *For any* Γ *and* α*,*

$$\Gamma \vdash_{\mathcal{N}} \alpha \iff \Gamma \vdash_a \alpha.$$

The Algebraic Completeness Theorem provides us with a canonical Heyting algebra $\mathcal{A}_\emptyset$, consisting of the equivalence classes of formulas under the equivalence relation induced by intuitionistic provable equivalence.

The Algebraic Soundness Theorem shows that any function from the propositional letters to a Heyting algebra $\mathcal{A}$, i.e. any environment on $\mathcal{A}$, can be extended to a homomorphism of Heyting algebras from $\mathcal{A}_\emptyset$ to $\mathcal{A}$, i.e. to the canonical interpretation associated with the environment. This property is concisely expressed by saying that $\mathcal{A}_\emptyset$ is the *free Heyting algebra on countably many generators*. More precisely, the generators are the equivalence classes of propositional letters, which are countably many because distinct letters cannot be provably equivalent.

Exercise 18.4.3 a) $\mathcal{A}_\emptyset$ *is not a topological Heyting algebra*. (Hint: it is enough to show that $\mathcal{A}_\emptyset$ has no point. If α corresponds to a point, then $[\![\alpha]\!] \neq 1$ and α is unprovable. Choose any letter p not occurring in it. Then $[\![p]\!] \sqcap [\![\neg p]\!] = 0 \sqsubseteq [\![\alpha]\!]$, but neither $[\![p]\!] \sqsubseteq [\![\alpha]\!]$ nor $[\![\neg p]\!] \sqsubseteq [\![\alpha]\!]$, otherwise $p \to \alpha$ or $\neg p \to \alpha$ would be provable, which is impossible.)

From the Algebraic Soundness Theorem we obtain as usual the following result, which provides a first characterization of Heyting algebras.

Theorem 18.4.4 First Representation for Heyting Algebras (Tarski [1935]) *Any Heyting algebra is isomorphic to a Lindenbaum algebra for the Intuitionistic Propositional Calculus.*

Proof. As in 5.4.1. □

Compact open sets of Stone topologies

The Stone Theorem provides a second characterization of Heyting algebras.

Theorem 18.4.5 Second Representation for Heyting Algebras (Stone [1937], McKinsey and Tarski [1946]) *Any Heyting algebra is isomorphic to a subalgebra of the algebra of open sets of its Stone space.*

Proof. See 5.4.2, 17.1.28 and 17.2.15. □

Our next goal is to characterize in a topological way the relevant subalgebras of Stone spaces. The following turns out to be the crucial notion.

Definition 18.4.6 *A subset X of a topological space is* **compact** *if, whenever it is covered by an union of open sets, it is already covered by a finite subunion.*

To prove the next result we show that the Stone topologies are actually generated by their compact open sets.

Theorem 18.4.7 Third Representation for Heyting Algebras (Stone [1937a]) *Any Heyting algebra is isomorphic to the algebra of compact open sets of its Stone topology.*

Proof. The proof of 17.1.28 shows that if $\mathcal{F}_A^p$ is the set of all prime filters on A, then the function $f : A \to \mathcal{P}(\mathcal{F}_A^p)$ defined as follows:

$$f(x) = \text{the set of all prime filters containing } x$$

is a one-one homomorphism of Heyting algebras.

It thus only remains to characterize $f(A)$ as the set of compact open sets.

- *every compact open set is in $f(A)$*
  Let X be a compact open set. Since X is open and $f(A)$ generates the Stone topology, there is a subset B of A such that
  $$X = \bigcup_{x \in B} f(x).$$
  Since X is compact, there is a finite subset $\{x_1, \ldots, x_n\}$ of B such that
  $$X = f(x_1) \cup \cdots \cup f(x_n).$$
  Then
  $$X = f(x_1 \sqcup \cdots \sqcup x_n)$$
  because f preserves $\sqcup$, and thus $X \in f(A)$.

- *every element of $f(A)$ is compact open*
  We first prove by contradiction that $f(1)$, i.e. the whole space $\mathcal{F}^p_A$, is compact. Suppose
  $$f(1) = \bigcup_{x \in B} f(x)$$
  but, for every finite subset $\{x_1, \ldots, x_n\}$ of B,
  $$f(1) \neq f(x_1) \cup \cdots \cup f(x_n).$$
  Then
  $$f(1) \neq f(x_1 \sqcup \cdots \sqcup x_n)$$
  because f preserves $\sqcup$, and
  $$1 \neq x_1 \sqcup \cdots \sqcup x_n$$
  by one-onenness of f.

  We want to find a prime filter F containing no $x \in B$, contradicting the fact that
  $$\mathcal{F}^p_A = f(1) = \bigcup_{x \in B} f(x),$$
  i.e. that every prime filter contains some x for $x \in B$.

  To find F it is enough to find a prime ideal I containing every $x \in B$, and then let F be its complement (since, on any lattice, I is a prime ideal if and only if its complement is a prime filter[7]). Consider then the ideal generated by B, which (as in 5.1.12.b) consists of the downward closure of the set of all finite joins of elements of B. Such an ideal is proper because, as noted above, all finite joins of elements of B are $\neq 1$.

  Then the set of all proper ideals containing B is non empty and partially ordered by inclusion, and every nonempty chain has a l.u.b. (which is just the union of the chain). By Zorn's Lemma, there is a maximal ideal I containing B, and such an ideal is prime as in the (dual) proof of 17.1.28.

[7] We show, for example, that if I is a prime ideal, then $\overline{I}$ is a filter (which is what is needed above):

- *if $x \in \overline{I}$ and $x \sqsubseteq y$, then $y \in \overline{I}$*
  Suppose $y \in I$. By downward closure of I, $x \in I$.
- *if $x, y \in \overline{I}$, then $x \sqcap y \in \overline{I}$*
  Suppose $x \sqcap y \in I$. By primality of I, $x \in I$ or $y \in I$.
- *if $x \sqcup y \in \overline{I}$, then $x \in \overline{I}$ or $y \in \overline{I}$*
  Suppose $x, y \in I$. By closure under $\sqcup$ of I, $x \sqcup y \in I$.

The proof that $f(a)$ is compact is a variation of the one just given for $f(1)$. Indeed, suppose

$$f(a) \subseteq \bigcup_{x \in B} f(x)$$

but, for every finite subset $\{x_1, \ldots, x_n\}$ of B,

$$f(a) \not\subseteq f(x_1) \cup \cdots \cup f(x_n).$$

Then

$$f(a) \not\subseteq f(x_1 \sqcup \cdots \sqcup x_n)$$

because f preserves $\sqcup$, and

$$a \not\sqsubseteq x_1 \sqcup \cdots \sqcup x_n$$

because f preserves $\sqsubseteq$, in particular

$$1 \neq x_1 \sqcup \cdots \sqcup x_n.$$

As above, there is a prime filter F containing a but no element of B (since the ideal generated by B does not contain a), contradiction. □

Corollary 18.4.8 (Stone [1937]) *Given an arbitrary Heyting algebra, its (dual) Stone space is compact, and its (dual) Stone topology is generated by the compact open sets.*

Proof. The proof of 18.4.7 proves the assertion for the Stone space $\mathcal{F}^p_A = f(1)$ of a Heyting algebra A. A dual proof works for the dual Stone space. □

18.5 Algebraic Heyting Algebras ⋆

In the present section we introduce an abstract version of the notion of compactness, as well as of the property of (dual) Stone topologies of being generated by a Heyting algebra of compact open sets.

Definition 18.5.1 (Birkhoff and Frink [1948], Nachbin [1949]) *Given a complete lattice A, an element a is called* **compact** *if, whenever $a \sqsubseteq \bigsqcup X$, there is a finite subset u of X such that $a \sqsubseteq \bigsqcup u$.*

A lattice A is called **algebraic** *if it is complete, and every element is the l.u.b. of the compact elements below it.*

The set of compact elements of a complete lattice A is indicated by $\mathbf{K(A)}$.

As usual, we can say that a complete lattice A has **enough compact elements** if, for every x and y in A,

$$x \neq y \implies k(x) \neq k(y),$$

where

$$k(x) = \{a : a \text{ compact } \wedge \ a \sqsubseteq x\}.$$

In other words, if $x \neq y$, then there is a compact element below one of x and y but not below the other. Then *a complete lattice is algebraic if and only if it has enough compact elements*. However, since the proof of 18.5.4 does not make any use of the function k, it is simpler to confine to the definition above.

Exercises 18.5.2 a) *If the $\sqcap\bigsqcup$-distributive law holds, then a is compact if and only if, whenever $a = \bigsqcup X$, there is a finite subset u of X such that $a = \bigsqcup u$.* (Hint: suppose $a \sqsubseteq \bigsqcup X$. Then

$$a = a \sqcap (\bigsqcup X) = \bigsqcup\{a \sqcap x : x \in X\},$$

and there is a finite subset u of X such that

$$a = \bigsqcup\{a \sqcap x : x \in u\} = a \sqcap (\bigsqcup u),$$

i.e. $a \sqsubseteq \bigsqcup u$.)

b) *An element of a complete lattice is a strong co-point if and only if it is a compact co-point.* (Hint: let $a \sqsubseteq \bigsqcup X$. If a is compact, then $a \sqsubseteq \bigsqcup u$ for some finite $u \subseteq X$. And if a is a co-point, then $a \sqsubseteq x$ for some $x \in u$.)

c) *Every finite lattice is algebraic.* (Hint: in a finite lattice every element is compact.)

d) *A complete linear ordering is algebraic if and only if any two distinct elements are separated by a gap, i.e. by two elements with nothing in between.* (Hint: first notice that in a linear ordering an element is compact if and only if it is 0 or it has an immediate predecessor. Indeed, if x has an immediate predecessor y, then any ideal of elements $\sqsubset x$ has l.u.b. $\sqsubseteq y \sqsubset x$. If x has no immediate predecessor, then the set I of all $y \sqsubset x$ is an ideal with l.u.b. x, to which x does not belong.

If the ordering is algebraic and $x \sqsubset y$, then $y \not\sqsubseteq x$, and there is a compact element $k \sqsubseteq y$ such that $k \not\sqsubseteq x$. Then $x \sqsubset k$, and hence k and its predecessor are between x and y. Conversely, if the distinct elements are separated by a gap and $x \neq 0$, then x is the l.u.b. of elements with an immediate predecessor.)

e) *Every power set is algebraic.* (Hint: in a power set an element is compact if and only if it is finite, because every set is the union of its finite subsets.)

f) *Every algebraic c.p.o. (see 6.3.28) is an algebraic Heyting algebra.* (Hint: an algebraic c.p.o. is a complete lattice with enough strong co-points.)

It follows from 18.5.2.c that not every algebraic lattice is distributive, and in particular *not every algebraic lattice is a Heyting algebra*. On the other hand, lack of distributivity is the only obstruction, since *a distributive algebraic lattice is a*

Heyting algebra (by 18.6.4 below, an algebraic lattice is continuous; and as noticed after 18.6.3, a distributive continuous lattice is a Heyting algebra).

The next result, together with 18.6.4 and 18.6.5, locates the class of algebraic Heyting algebras in the spectrum of Heyting algebras dealt with in the present section.

Proposition 18.5.3 *Every complete Heyting algebra with enough strong co-points is algebraic, but there are algebraic Heyting algebras without (strong) co-points.*

Proof. A complete Heyting algebra with enough strong co-points is algebraic because, by 18.5.2.b, a strong co-point is compact. Then, if every element is the l.u.b. of the strong co-points below it, it is also the l.u.b. of the compact elements below it.

An example of an algebraic Heyting algebra without (strong) co-points is the *algebra of open sets of the Cantor space* 2^ω, i.e. the set of all 0,1-valued functions on ω, with the topology generated by the sets $\{f : f \supseteq \sigma\}$ of functions having a common fixed initial segment σ. The Cantor space is algebraic because such sets are compact, by König's Lemma, and generate the topology by definition. And there is no co-point because each such set is the union of different open sets, since each function having σ as an initial segment must have either 0 or 1 as its next value:

$$\{f : f \supseteq \sigma\} = \{f : f \supseteq \sigma * 0\} \cup \{f : \supseteq \sigma * 1\}. \quad \Box$$

Topologies generated by their compact open sets

The next result provides a topological characterization of algebraic Heyting algebras.

Theorem 18.5.4 Topological Characterization of Algebraic Heyting Algebras (Hofmann and Keimel [1972]) *The algebraic Heyting algebras are, up to isomorphism, exactly the algebras of open sets of topologies generated by their compact open sets.*

Proof. Sufficiency is immediate by Definition 18.5.1, since the algebraic notion of compactness is patterned on the topological one.

To show necessity, we prove in 18.6.4 and 18.6.5 below that an algebraic Heyting algebra A has enough points. By 18.3.4, such an algebra is isomorphic to the algebra of open sets of the form

$$p(x) = \{a : a \text{ point } \wedge\ a \not\supseteq x\}.$$

It is thus enough to notice that such a topology is generated by the compact open sets. Since the topology is generated by $p(A)$, it is enough to show that every element of $p(A)$ is the union of the compact open sets contained in it. Since every element of A is the l.u.b. of the compact elements below it, and p preserves arbitrary l.u.b.'s, it is enough to prove that the image of a compact element is compact.

Let thus a be compact on A, and $p(a) \subseteq \bigcup_{x \in X} p(x)$. Then $p(a) \subseteq p(\bigsqcup X)$ because p preserves arbitrary l.u.b.'s, and then $a \sqsubseteq \bigsqcup X$ because p is an isomorphism. But a is compact, so $a \sqsubseteq \bigsqcup u$ for some finite $u \subseteq X$. Then

$$p(a) \subseteq p(\bigsqcup u) = \bigcup_{x \in u} p(x),$$

i.e. $p(a)$ is compact. $\square$

Exercises 18.5.5 Arithmetic Heyting algebras. A Heyting algebra A is called **arithmetic** if it is algebraic, and the compact elements are a subalgebra of A, i.e. they are closed under the Heyting algebra operations $\sqcap$, $\sqcup$ and $\Rightarrow$, and contain 0 and 1.

a) *The set of compact elements of a complete lattice A contains 0 and is closed under* $\sqcup$. (Hint: if a and b are compact, suppose $(a \sqcup b) \sqsubseteq \bigsqcup X$. Then there are finite subsets u_a and u_b of X such that $a \sqsubseteq \bigsqcup u_a$ and $b \sqsubseteq \bigsqcup u_b$, i.e. $(a \sqcup b) \sqsubseteq \bigsqcup (u_a \cup u_b)$.)

b) *Every finite Heyting algebra is arithmetic.* (Hint: in a finite lattice every element is compact.)

c) *An algebraic linear ordering is arithmetic if and only if 1 has an immediate predecessor.* (Hint: since in a linear ordering $x \sqcap y = x$ and $x \Rightarrow y$ is either y or 1, the compact elements are automatically a subalgebra if 1 is compact.)

d) *The only arithmetic power sets are the finite ones.* (Hint: closure under $\Rightarrow$ implies that if a is finite then so is $\overline{a} = a \Rightarrow 0$, i.e. a must be cofinite too.)

e) *Every arithmetic Heyting algebra is algebraic, but there are algebraic Heyting algebras that are not arithmetic.* (Hint: by 18.5.2.e and part d).)

f) *The arithmetic Heyting algebras are, up to isomorphism, exactly the algebras of open sets of compact topologies such that the compact open sets generate the topology and form a subalgebra.* (Hint: by 18.5.4 and part a).)

Exercises 18.5.6 Compact elements in Heyting algebras of ideals. Given a lattice A, let $\mathcal{I}_A$ be the set of its ideals.

a) *If A is a Heyting algebra, then $\mathcal{I}_A$ is an arithmetic Heyting algebra.* (Hint: we first prove that the compact elements of $\mathcal{I}_A$ are exactly the principal ideals $\downarrow a$ generated by elements $a \in A$. Suppose I is a compact element of $\mathcal{I}_A$. Since I is the l.u.b. of the principal ideals $\downarrow a$ for $a \in I$, it is also the l.u.b. of a finite set of principal ideals $\downarrow a_1, \ldots, \downarrow a_n$, and hence it is the principal ideal $\downarrow (a_1 \sqcup \cdots \sqcup a_n)$. Conversely, suppose $(\downarrow a) \subseteq \bigsqcup_{x \in X} J_x$: then $a \in \bigsqcup_{x \in X} J_x$, i.e. $a \sqsubseteq a_1 \sqcup \cdots \sqcup a_n$ for some $a_i \in J_{x_i}$; thus $a \in J_{x_1} \sqcup \cdots \sqcup J_{x_n}$, i.e. $(\downarrow a) \subseteq J_{x_1} \sqcup \cdots \sqcup J_{x_n}$, and $\downarrow a$ is compact.

That $\mathcal{I}_A$ is arithmetic now follows from the facts that every ideal is the l.u.b. of the principal ideals generated by its elements, and that the principal ideals form a subalgebra of $\mathcal{I}_A$, i.e. $(\downarrow a) \cap (\downarrow b) = \downarrow (a \sqcap b)$, $(\downarrow a) \Rightarrow (\downarrow b) = \downarrow (a \Rightarrow b)$, and $(\downarrow 1) = A$, since preservation of $\sqcup$ and 0 is automatic by 18.5.5.a.

For example, $(\downarrow a) \Rightarrow (\downarrow b)$ is the smallest ideal containing all ideals I such that $I \cap (\downarrow a) \subseteq (\downarrow b)$. If x belongs to it, then $x \sqsubseteq x_1 \sqcup \cdots \sqcup x_n$, with $x_i \sqcap a \sqsubseteq b$. Thus $x \sqcap a \sqsubseteq b$, i.e. $x \sqsubseteq (a \Rightarrow b)$, and $x \in \downarrow (a \Rightarrow b)$. Conversely, $\downarrow (a \Rightarrow b)$ is an ideal of elements x such that $x \sqcap (a \Rightarrow b)$, i.e. $x \sqcap a \sqsubseteq b$, and thus an ideal contained in $(\downarrow a) \Rightarrow (\downarrow b)$.)

b) *Any arithmetic Heyting algebra is isomorphic to $\mathcal{I}_A$, for some Heyting algebra A.* (Hint: given an arithmetic Heyting algebra B, let $K(B)$ be the set of its compact elements. Since B is arithmetic, $K(B)$ is a subalgebra of A, and hence a Heyting algebra. It is thus enough to show that $B \simeq \mathcal{I}_{K(B)}$, via the function

$$k(a) = \{c : c \text{ is compact} \wedge c \sqsubseteq a\}.$$

Indeed, $k(a)$ is an ideal of compact elements of B because it is trivially closed downward, and the compact elements are closed under $\sqcap$. Moreover, k is onto because each ideal of compact elements is the image of the l.u.b. of its elements. And k is one-one because the compact elements generate B, i.e. every element of B is the l.u.b. of the compact elements below it, and if two elements are distinct there must be a compact element below one but not the other.)

Exercise 18.5.7 Coherent spaces. (Stone [1937], Serre [1955]) A topological space is called **coherent** if:

1. as a topology, it is sober (see 18.3.5)
2. as a Heyting algebra, it is arithmetic (see 18.5.1).

A topology is homeomorphic to the dual Stone topology of a Heyting algebra if and only if it is coherent. (Hint: given a Heyting algebra A, by 18.3.6.b the dual Stone topologies of A and of $\mathcal{I}_A$ are the same. By 18.3.6.a, $\mathcal{I}_A$ is a complete Heyting algebra with enough points. By 18.3.5, its dual Stone topology is thus sober, and it satisfies 1. A proof dual to that of 18.4.7 shows that it also satisfies 2.

Conversely, let $\Omega(X)$ be a coherent space, and $K\Omega(X)$ be the set of its compact open sets. By 1 and 18.3.5.d,

$$\Omega(X) \simeq \Omega(\mathrm{Pt}(\Omega(X))),$$

By 2 and 18.5.6.b,

$$\Omega(X) \simeq \mathcal{I}_{K\Omega(X)},$$

i.e.

$$\Omega(\mathrm{Pt}(\Omega(X))) \simeq \Omega(\mathrm{Pt}(\mathcal{I}_{K\Omega(X)})).$$

Thus $\Omega(X)$ is homeomorphic to the dual Stone topology of the Heyting algebra $\mathcal{I}_{K\Omega(X)}$ (as a Heyting algebra with enough points), and by 18.3.6.b also to the dual Stone topology of the Heyting algebra $K\Omega(X)$ (as an arbitrary Heyting algebra).)

18.6 Continuous Heyting Algebras ⋆

Strong co-points and compact elements can be considered as approximations to elements, generating the lattice when there are enough of them. We now relativize the notion of compactness and introduce a further notion of approximation, which turns out to be the most appropriate for applications.

Definition 18.6.1 (Scott [1972]) *An element a of a lattice is* **way below** *another element x ($\boldsymbol{a \ll x}$) if a is in any ideal I such that $x \sqsubseteq \bigsqcup I$.*

A lattice is **continuous** *if it is complete and, for all x,*

$$x = \bigsqcup \{a : a \ll x\}.$$

If we let

$$i(x) = \bigcap \{I : I \text{ ideal } \wedge \ x \sqsubseteq \bigsqcup I\}$$

and

$$d(x) = \bigcap \{D : D \text{ downward closed } \wedge \ x \sqsubseteq \bigsqcup D\},$$

then

$$d(x) \subseteq i(x) = \{a : a \ll x\} \subseteq \{a : a \sqsubseteq x\},$$

where the first inclusion follows from the fact that every ideal is downward closed, and the second from the fact that $\{a : a \sqsubseteq x\}$ is an ideal with l.u.b. x. In particular, $\ll$ *is stronger than* $\sqsubseteq$, i.e.

$$a \ll x \implies a \sqsubseteq x.$$

Proposition 18.6.2 $\ll$ *is transitive and dense.*

Proof. Transitivity is immediate. Suppose $a \ll b \ll c$ and $c \sqsubseteq \bigsqcup I$. Then $b \in I$, because $b \ll c$. And $a \in I$, because I is an ideal and $a \sqsubseteq b$ (since $a \ll b$).

To prove density, suppose $a \ll b$ and consider the set

$$I = \{x : (\exists y)(x \ll y \ll b)\}.$$

It is enough to show that I is an ideal such that $b \sqsubseteq \bigsqcup I$. Then $a \in I$, because $a \ll b$. And $a \ll y \ll b$ for some y, by definition of I.

- *I is downward closed*
  Suppose $z \sqsubseteq x$ and $x \in I$. Then $x \ll y \ll b$ for some y, by definition of I. And $z \ll y$ because $z \sqsubseteq x \ll y$, by definition of $\ll$. Thus $z \ll y \ll b$, and $z \in I$.

- *I is closed under $\sqcup$*
  Suppose $x_1 \ll y_1 \ll b$ and $x_2 \ll y_2 \ll b$. Since $\{z : z \ll b\}$ is an ideal, being an intersection of ideals, $y_1 \sqcup y_2 \ll b$. Since $\{z : z \ll y_1 \sqcup y_2\}$ is an ideal, $x_1 \sqcup x_2 \ll y_1 \sqcup y_2$. Thus $x_1 \sqcup x_2 \ll y_1 \sqcup y_2 \ll b$, and $x_1 \sqcup x_2 \in I$.

- $b \sqsubseteq \bigsqcup I$
  Since I is an ideal of elements $\ll b$, $\bigsqcup I \sqsubseteq b$. Suppose $\bigsqcup I \sqsubset b$. Since the elements way below b have l.u.b. b, there is $y \ll b$ such that $y \not\sqsubseteq \bigsqcup I$. Since the elements way below y have l.u.b. y, there is $x \ll y$ such that $x \not\sqsubseteq \bigsqcup I$. Thus $x \ll y \ll b$, i.e. $x \in I$, and $x \not\sqsubseteq \bigsqcup I$, contradiction. □

Exercises 18.6.3 a) *An element a of a complete lattice is compact if and only if $a \ll a$.* (Hint: one direction follows from the fact that ideals are closed under l.u.b.'s of finite sets. The other direction follows by considering the ideal generated by X, i.e. the downward closure of the set of l.u.b.'s of finite subsets of X.)

b) *In a linear ordering, $a \ll x$ if and only if $a \sqsubset x$ or $a = x \ll x$. And $x \ll x$ if and only if $x = 0$ or x has an immediate predecessor.* (Hint: see 18.5.2.d.)

c) *In a finite lattice, $a \ll x$ if and only if $a \sqsubseteq x$. And $x \ll x$ for every x.*(Hint: in a finite lattice $\bigsqcup I \in I$. So, if $x \sqsubseteq \bigsqcup I$, then $x \in I$. And if $a \sqsubseteq x$, then $a \in I$ too.)

d) *In a $\sqcap\bigsqcup$-distributive lattice, $a \ll x$ if and only if a is in any ideal I such that $x = \bigsqcup I$.* (Hint: suppose $x \sqsubseteq \bigsqcup I$. By $\sqcap\bigsqcup$-distributivity

$$x = x \sqcap \bigsqcup I = \bigsqcup \{x \sqcap i : i \in I\},$$

and the right-hand-side is the l.u.b. of an ideal. So $a \sqsubseteq x \sqcap i$ for some i, and $a \in I$.)

e) *In a power set, $a \ll x$ if and only if a is a finite subset of x. And $x \ll x$ if and only if x is finite.* (Hint: let I be the the ideal of finite subsets of x, so that $x = \bigcup I$. If $a \ll x$, then $a \in I$.)

f) *Every finite lattice is continuous.* (Hint: by part c).)

g) *Every complete linear ordering is continuous.* (Hint: by part b).)

h) *Every power set is continuous.* (Hint: by part e).)

It follows from 18.6.3.f that not every continuous lattice is distributive, and in particular *not every continuous lattice is a Heyting algebra*. On the other hand, lack of distributivity is the only obstruction, since *a distributive continuous lattice is a Heyting algebra* (as noticed after the proof of 18.6.5 below, a distributive continuous lattice has enough points; and as noticed after the proof of 18.3.4, a complete lattice with enough points is a Heyting algebra).

The next two results locate the class of continuous Heyting algebras in the spectrum of Heyting algebras dealt with in the present chapter.

Proposition 18.6.4 *Every algebraic lattice is continuous, but there are continuous Heyting algebras that are not algebraic.*

Proof. To prove the first part, it is enough to notice that, by 18.6.3.a, if a is a compact element such that $a \sqsubseteq x$, then $a \ll a \sqsubseteq x$, and hence $a \ll x$. Then, if every element is the l.u.b. of the compact elements below it (i.e. if the lattice is algebraic), it is also the l.u.b. of the elements way below it (i.e. the lattice is continuous).

To prove the second part, we want to exhibit a continuous Heyting algebra that is not algebraic. The *algebra of open sets of the euclidean space $I\!R$* provides an example.

First, such a Heyting algebra is not algebraic because the only compact open set is $\emptyset$.

Second, we prove that if V is an open set and U is an open interval whose closure $cl(U)$ is contained in V, then $U \ll V$. Continuity then follows from the fact that the open intervals generate the topology (more specifically, that every open set is a union of open intervals).

Suppose $V \subseteq \bigcup_{i \in I} A_i$, where $\mathcal{I} = \{A_i\}_{i \in I}$ is an ideal of open sets. Then

$$U \subseteq cl(U) \subseteq V \subseteq \bigcup_{i \in I} A_i.$$

Since $cl(U)$ is a closed interval, it is compact. Then there are $i_1, \ldots, i_n$ such that

$$U \subseteq cl(U) \subseteq A_{i_1} \cup \cdots \cup A_{i_n}.$$

But $A_{i_1} \cup \cdots \cup A_{i_n} \in \mathcal{I}$ by closure under finite $\cup$, and then $U \in \mathcal{I}$ by downward closure of $\mathcal{I}$. □

Proposition 18.6.5 (Papert [1959]) *Every continuous Heyting algebra has enough points, but there are complete Heyting algebras with enough points that are not continuous.*

Proof. To prove the first part, we refer to the proof of 18.3.3. To be able to extend its last part, given $x \not\sqsubseteq y$ we need to find a filter F with the properties used there.

Since the elements way below x have l.u.b. x, there is $a \ll x$ such that $a \not\sqsubseteq y$. By density of $\ll$ (18.6.2), there is an infinite descending chain

$$a \ll \cdots \ll x_2 \ll x_1 \ll x.$$

We consider the set

$$F = \{z : (\exists n)(x_n \sqsubseteq z)\},$$

and notice the following properties:

- F is a filter, being the union of the principal filters generated by the x_n's. In particular, F is closed upward and under $\sqcap$.
- F contains only elements above a. In particular $y \notin F$.
- No element of $\overline{F}$ is above x. This follows from the fact that $x \in F$ (because $x_1 \ll x$, and hence $x_1 \sqsubseteq x$), since F is upward closed.
- Every maximal element in $\overline{F}$ is $\sqcap$-irreducible, and hence a point (by 18.3.2.c). This follows from the fact that F is closed under $\sqcap$.
- If I is an ideal contained in $\overline{F}$, then $\bigsqcup I$ is also in $\overline{F}$.[8] Otherwise $\bigsqcup I$ is in F, and there is some n such that $x_n \sqsubseteq \bigsqcup I$. Since $x_{n+1} \ll x_n$, x_{n+1} must be below some element of I, which is impossible because I is contained in $\overline{F}$, and F is upward closed.

[8] This property of F (of being inaccessible by l.u.b.'s of ideals in $\overline{F}$) is characteristic of open sets in the Scott topology, see 18.6.14.

Let now C be any maximal chain in $\overline{F}$ containing y, which exists by Zorn's Lemma. The downward closure of C is an ideal I contained in $\overline{F}$, and thus $\bigsqcup I$ is also in $\overline{F}$. Since C is maximal, $\bigsqcup I$ is a maximal element of $\overline{F}$, and hence a point above y (by the choice of C), but not above in x (because it is in $\overline{F}$).

To prove the second part, we want to exhibit a complete Heyting algebra with enough points that is not continuous. An example is provided by the *algebra of open sets of the Baire space* ω^ω, i.e. the set of all functions on ω, with the topology generated by the sets $\{f : f \supseteq \sigma\}$ of functions having a common fixed initial segment σ.

First, we prove that such an algebra has enough points. Indeed, for any function f, the set $\overline{\{f\}}$ is a point: it is open because it is the union of the basic open sets defined by initial segments σ differing from f on at least one point; and it is a point because if $\overline{\{f\}} \supseteq A \cap B$, then at least one of A and B does not contain f, and it is thus contained in $\overline{\{f\}}$. Moreover, there are enough points because if A and B are distinct open sets, then there is a function f in one but not in the other, say $f \in A - B$. Then the point $\overline{\{f\}}$ contains B, but not A.

Second, to prove that such an algebra is not continuous it is enough to show that if A and B are open sets such that $A \ll B$, then $A = \emptyset$. Suppose $A \neq \emptyset$: it is enough to show that it is possible to decompose the whole space ω^ω into infinitely many disjoint open sets, each containing at least an element of A. If I is the ideal generated by such open sets, then $B \subseteq \bigcup I$: since $A \ll B$, A should be in I, and hence be contained in a finite union of such open sets, contradiction.

If $A \neq \emptyset$, then A contains a basic open set, defined by an initial segment σ. On the one hand, the complement of such an open set is open, being the union of the basic open sets defined by initial segments incompatible with σ. On the other hand, such an open set is the disjoint union of infinitely many basic open sets, i.e. the ones defined by one element extensions of σ. The needed decomposition of the whole space is easily obtained from these open sets. □

Notice that the first part of the proof shows in particular that *a distributive continuous lattice has enough points*, since no use of $\Rightarrow$ was made in it. Distributivity, which was used when claiming that a $\sqcap$-irreducibile element is a point, is essential because the previous proof can be combined with that of 18.3.4 to show that *a distributive continuous lattice is a Heyting algebra*, while not every continuous lattice is such (by 18.6.3.f).

As for the second part of the proof, it will be put into a broader perspective by 18.6.11.c. In particular, the fact that the algebra of open sets of the Baire space is not continuous follows from the fact that the topology is T_2 but not locally compact (more precisely, the interior of any compact set is empty).

Locally quasi-compact topologies

We now look for a topological characterization of the class of continuous Heyting algebras.

By 18.4.8 the (dual) Stone topology of an arbitrary Heyting algebra is generated by compact open sets, since each $f(x)$ is compact. We now look at the dual Stone topology of a continuous Heyting algebra (as an algebra with enough points), and show that it is generated by the interiors of compact open sets, in the following sense.

Definition 18.6.6 *A topological space is called* **locally quasi-compact** *if, for any element x and any open set V such that $x \in V$, there are a compact set C_x and an open set O_x such that*

$$x \in O_x \subseteq C_x \subseteq V.$$

A topological space is called **locally compact** *if it is locally quasi-compact and T_2.*

A typical example of a locally (quasi-)compact, but not compact space is given by $I\!\!R$ with the usual topology. The next proof generalizes the fact, proved in 18.6.4, that the algebra of open sets of $I\!\!R$ is continuous.

Proposition 18.6.7 (Day and Kelly [1970]) *Every algebra of open sets of a locally quasi-compact topology is a continuous Heyting algebra.*

Proof. Given an open set V, for any $x \in V$ we consider the compact set C_x and the open set O_x provided by the definition of local quasi-compactness. Since $V = \bigcup_{x \in V} O_x$, to prove that the algebra of open sets is continuous it is enough to show that $O_x \ll V$. Then each open set is the l.u.b. of open sets way below it.

Suppose $V \subseteq \bigcup_{i \in I} A_i$, where $\mathcal{I} = \{A_i\}_{i \in I}$ is an ideal of open sets. Then

$$O_x \subseteq C_x \subseteq V \subseteq \bigcup_{i \in I} A_i.$$

By definition, from any family of open sets whose union covers a compact set, we can extract a finite subfamily with the same property. Then, by compactness of C_x,

$$O_x \subseteq C_x \subseteq A_{i_1} \cup \cdots \cup A_{i_n}$$

for some $i_1, \ldots, i_n$. But $A_{i_1} \cup \cdots \cup A_{i_n} \in \mathcal{I}$ by closure under finite $\cup$, and then $O_x \in \mathcal{I}$ by downward closure of $\mathcal{I}$. □

The next exercise clarifies in which sense $\ll$ can be considered as a notion of relative compactness.

Exercise 18.6.8 *In the algebra of open sets of a locally quasi-compact topology, $U \ll V$ if and only if U is an open set contained in a compact set contained in V.* (Hofmann and Lawson [1978]) (Hint: one direction has just been proved. Conversely, suppose $U \ll V$. As above, $V = \bigcup_{x \in V} O_x$. Since $U \ll V$, there are $x_1, \ldots, x_n$ such that

$$U \subseteq O_{x_1} \cup \cdots \cup O_{x_n} \subseteq C_{x_1} \cup \cdots \cup C_{x_n} \subseteq V,$$

and $C_{x_1} \cup \cdots \cup C_{x_n}$ is compact.)

We turn now to the converse of the previous result.

Proposition 18.6.9 (Hofmann and Lawson [1978]) *A continuous Heyting algebra is isomorphic to the algebra of open sets of a locally quasi-compact topology.*

Proof. By 18.6.5 a continuous Heyting algebra has enough points, and by 18.3.4 it is isomorphic to the algebra of open sets of the form

$$p(x) = \{a : a \text{ point } \wedge \ a \not\sqsupseteq x\}.$$

It is thus enough to show that such a topology is locally quasi-compact, i.e. that for any element y and any open set $p(x)$ such that $y \in p(x)$, there is a compact set C_y and an open set O_y such that

$$y \in O_y \subseteq C_y \subseteq p(x).$$

Since $y \in p(x)$, y is a point such that $x \not\sqsubseteq y$. As in the proof of 18.6.5, there is $a \ll x$ such that $a \not\sqsubseteq y$, and we can find a filter F with the following properties:

- F contains only elements $\sqsupseteq a$.
- No element of $\overline{F}$ is above x.
- If I is an ideal contained in $\overline{F}$, then $\bigsqcup I$ is also in $\overline{F}$.

Since y is a point and $y \not\sqsupseteq a$, $y \in p(a)$. But $p(a)$ is open, so we can let $O_y = p(a)$. Moreover, $p(a)$ is a set of points not above a, and hence in $\overline{F}$ and not above x. If we let C_y be the set of points in $\overline{F}$, we then automatically have

$$y \in O_y \subseteq C_y \subseteq p(x).$$

It only remains to show that C_y is compact, and we prove this by contrapositive.

First we notice that the downward closure of C_y coincides with $\overline{F}$. Indeed, if an element is below a point in $\overline{F}$, then it cannot be in F, because F is upward closed (being a filter), and hence its complement is downward closed. Conversely, any element of $\overline{F}$ is bounded by a point in $\overline{F}$, by the proof of 18.6.5.

Given now a subset B of A we suppose that, for any finite subset $\{x_1, \ldots, x_n\}$ of B,

$$C_y \not\subseteq p(x_1) \cup \cdots \cup p(x_n).$$

Since p preserves $\sqcup$,

$$C_y \not\subseteq p(x_1 \sqcup \cdots \sqcup x_n).$$

This means that there is a point in $\overline{F}$ above $x_1 \sqcup \cdots \sqcup x_n$. Since $\overline{F}$ is the downward closure of C_y, this means that the ideal I generated by B is contained in $\overline{F}$. By the properties of F, then $\bigsqcup I$ is in $\overline{F}$, and hence so is $\bigsqcup B$. Again because $\overline{F}$ is the downward closure of C_y, this means that $\bigsqcup B$ is bounded by an element in C_y, i.e. by a point in $\overline{F}$. Then

$$C_y \not\subseteq p(\bigsqcup B).$$

Since p preserves $\bigsqcup$,

$$C_y \not\subseteq \bigcup_{x \in B} p(x). \quad \square$$

We can now put the two halves together.

Theorem 18.6.10 Topological Characterization of Continuous Heyting Algebras (Day and Kelly [1970], Hofmann and Lawson [1978]) *The continuous Heyting algebras are, up to isomorphism, exactly the algebras of open sets of locally quasi-compact topologies.*

Proof. By 18.6.7 and 18.6.9. $\square$

By describing topologies with reference only to their open sets, and not to the points of their underlying spaces, 18.3.4 allows us to see the theory of Heyting algebras as a kind of *pointless topology*, in which topological properties can be described in a purely algebraic way. 18.6.10 provides a paradigm in this direction, capturing the topological property of local quasi-compactness in terms of the algebraic property of continuity.

Exercises 18.6.11 a) *A topology is homeomorphic to the dual Stone topology of a continuous Heyting algebra if and only if it is sober and locally quasi-compact.* (Hint: by 18.3.5.d and 18.6.10.)

b) *A sober space is locally quasi-compact if and only if its algebra of open sets is continuous.* (Hint: the proof of 18.6.9 shows only that if the algebra of open sets of a topology is continuous, then the dual Stone space of the algebra is locally quasi-compact, and not that the given space is. But if the given space is sober, the dual Stone topology of its algebra of open sets is homeomorphic to it by 18.3.5.d.)

c) *A T_2 space is locally compact if and only if its algebra of open sets is continuous.* (Hint: by part b), since by 18.3.5.c a T_2 topology is sober.)

The next exercises introduce and discuss yet another interesting class of Heyting algebras.

Exercises 18.6.12 Completely distributive Heyting algebras. A lattice is $\bigwedge\bigsqcup$-**distributive** if, for any family $\{X_i\}_{i\in I}$ of subsets of A,

$$\bigwedge\{\bigsqcup X_i : i \in I\} = \bigsqcup\{\bigwedge\{f(i) : i \in I\} : f(i) \in X_i\},$$

where on the right-hand-side we consider all possible choice functions f for the family $\{X_i\}_{i\in I}$.

A lattice is $\bigsqcup\bigwedge$**-distributive** if it satisfies the dual condition.

A lattice is **completely distributive** if it is complete and both $\bigwedge\bigsqcup$-distributive and $\bigsqcup\bigwedge$-distributive.

a) *A completely distributive lattice is a Heyting algebra.* (Hint: by 5.3.5.)

b) *A complete lattice is* $\bigwedge\bigsqcup$*-distributive if and only if it is* $\bigsqcup\bigwedge$*-distributive.* (Raney [1952]) (Hint: it is enough to prove that $\bigwedge\bigsqcup$-distributivity implies $\bigsqcup\bigwedge$-distributivity, the converse implication being symmetric. It is also enough to prove

$$\bigwedge\{\bigsqcup\{f(i) : i \in I\} : f(i) \in X_i\} \sqsubseteq \bigsqcup\{\bigwedge X_i : i \in I\},$$

the converse inequality being automatic. We rewrite the left-hand-side in a form suitable for an application of $\bigwedge\bigsqcup$-distributivity, i.e. as $\bigwedge\{\bigsqcup Y_f : f \in F\}$, where

$$F = \{f : (\forall i)(f(i) \in X_i)\} \qquad \text{and} \qquad Y_f = \{f(i) : i \in I\}.$$

Then

$$\bigwedge\{\bigsqcup Y_f : f \in F\} = \bigsqcup\{\bigwedge\{g(f) : f \in F\} : g(f) \in Y_f\},$$

by $\bigwedge\bigsqcup$-distributivity.

It is now enough to show $(\forall g)(\exists i \in I)(X_i \subseteq \{g(f) : f \in F\})$. Suppose, for the sake of contradiction, that $(\exists g)(\forall i \in I)(X_i \not\subseteq \{g(f) : f \in F\})$. Then, for such a g, $(\forall i \in I)(\exists x_i \in X_i - \{g(f) : f \in F\})$. By letting $f(i) = x_i$ for all $i \in I$, we get an $f \in F$. Since g is a choice function for F, it should be $g(f) \in Y_f$, i.e. $g(f) = f(i) = x_i$ for some $i \in I$. But $x_i \notin \{g(f) : f \in F\}$, contradiction.)

c) *Any finite distributive lattice is completely distributive.* (Hint: on a finite lattice $\bigwedge$ and $\bigsqcup$ reduce to $\sqcap$ and $\sqcup$, and thus $\bigwedge\bigsqcup$-distributivity reduces to the usual distributivity.)

d) *Any complete linear ordering is completely distributive.* (Hint: to show $a \sqsubseteq b$, where $a = \bigwedge\{\bigsqcup X_i : i \in I\}$ and $b = \bigsqcup\{\bigwedge\{f(i) : i \in I\} : f(i) \in X_i\}$, suppose $b \sqsubset a$, and define a choice function f for $\{X_i\}_{i\in I}$, as follows: let $f(i)$ be an element of X_i strictly greater than b, which exists because by hypothesis $a \sqsubseteq \bigsqcup X_i$. This produces a contradiction both if there is no element between a and b (using the fact that then $f(i) \sqsupseteq a$), and if there is (by choosing $f(i)$ above it, and proceeding as in the first case).)

e) *Any power set is completely distributive.* (Hint: to show $a \subseteq b$, where $a = \bigcap\{\bigcup X_i : i \in I\}$ and $b = \bigcup\{\bigcap\{f(i) : i \in I\} : f(i) \in X_i\}$, let $x \in a$, and define a choice function f for $\{X_i\}_{i\in I}$, as follows. Let $f(i)$ be an element of X_i to which x belongs, which exists because by hypothesis $x \in \bigcup X_i$ for every $i \in I$. Then $x \in \bigcap\{f(i) : i \in I\}$, and so $x \in b$.)[9]

[9] The statement of part e) is equivalent to the Axiom of Choice (Collins [1954], Linton and Mikkelsen [1981]). Thus the use of the Axiom of Choice made in its proof, when defining f, is not avoidable.

f) *Every complete Heyting algebra with enough strong co-points is completely distributive, but there are completely distributive Heyting algebras without strong co-points.* (Büchi [1952], Raney [1952]) (Hint: to show $a \sqsubseteq b$, where $a = \bigwedge\{\bigsqcup X_i : i \in I\}$ and $b = \bigsqcup\{\bigwedge\{f(i) : i \in I\} : f(i) \in X_i\}$, let x be a strong co-point such that $x \sqsubseteq a$, and define a choice function f for $\{X_i\}_{i \in I}$, as follows: let $f(i)$ be an element of X_i greater than or equal to x, which exists because by hypothesis $x \sqsubseteq \bigsqcup X_i$ for every $i \in I$, and x is a strong co-point. Then $x \sqsubseteq \bigwedge\{f(i) : i \in I\}$, and so $x \sqsubseteq b$.

A complete dense linear ordering, e.g. any closed interval of the reals, is completely distributive by part d), but has no strong co-points by 18.2.2.g.)

g) *Every completely distributive Heyting algebra is continuous, but there are continuous Heyting algebras that are not completely distributive.* (Papert [1959]) (Hint: to prove the first part, it is enough to show that if A is completely distributive, then $x \sqsubseteq \bigsqcup d(x)$, where

$$d(x) = \bigcap\{D : D \text{ downward closed } \wedge \ x \sqsubseteq \bigsqcup D\}.$$

Consider the family $\{X_i\}_{i \in I}$ of all downward closed sets X_i such that $x \sqsubseteq \bigsqcup X_i$. Then, by complete distributivity, $x \sqsubseteq \bigwedge\{\bigsqcup X_i : i \in I\} = \bigsqcup\{\bigwedge\{f(i) : i \in I\} : f(i) \in X_i\}$. But $(\bigwedge\{f(i) : i \in I\}) \in \bigcap_{i \in I} X_i$ for any f, because $f(i) \in X_i$, and X_i is downward closed. Then $(\bigwedge\{f(i) : i \in I\}) \in d(x)$ by definition of $d(x)$ and the choice of the X_i's, and thus

$$\bigsqcup\{\bigwedge\{f(i) : i \in I\} : f(i) \in X_i\} \sqsubseteq \bigsqcup d(x).$$

To prove the second part, it is enough to notice that the algebra of open sets of the euclidean space $\mathbb{R}$ is continuous, but $\bigsqcup\bigwedge$-distributivity fails for the family $\{X_i\}_{i \in \mathbb{Z}}$, where X_i is the set of all open intervals containing the closed interval $[i, i+1]$.)

h) *There are completely distributive Heyting algebras that are not algebraic, as well as algebraic Heyting algebras that are not completely distributive.* (Hint: for the first part, use the example in part f) and 18.5.2.d. For the second part, use the example in 18.5.3 and 18.6.13.a.)

Exercises 18.6.13 Continuous Heyting algebras with enough co-points.

a) *A completely distributive Heyting algebra has enough co-points.* (Papert [1959]) (Hint: only $\bigwedge\bigsqcup$-distributivity was used in 18.6.12.g to show that completely distributive lattices are continuous. By the same argument, using the fact that $\bigsqcup\bigwedge$-distributivity also holds by 18.6.12.a, the lattice with reverse order is continuous, and hence it has enough points, i.e. the original lattice has enough co-points.)

b) *A continuous Heyting algebra with enough co-points is completely distributive.* (Kamara [1978]) (Hint: as in 18.6.12.f we want to show that $a \sqsubseteq b$, where $a = \bigwedge\{\bigsqcup X_i : i \in I\}$ and $b = \bigsqcup\{\bigwedge\{f(i) : i \in I\} : f(i) \in X_i\}$. Let x be a co-point such that $x \ll a$, and let $f(i)$ be an element of X_i greater than or equal to x, which exists because by hypothesis $x \ll a \sqsubseteq \bigsqcup X_i$ for every $i \in I$. Then $x \sqsubseteq \bigwedge\{f(i) : i \in I\}$, and so $x \sqsubseteq b$.)

c) *The continuous Heyting algebras with enough co-points are exactly the completely distributive Heyting algebras.* (Hint: from parts a) and b).)

Digression: Scott topologies $\star$

In the proofs of 18.6.5 and 18.6.9 a crucial role was played by filters F such that:

- if I is an ideal contained in $\overline{F}$, then $\bigsqcup I$ is also in $\overline{F}$.

The notion is sufficiently interesting to deserve special consideration.

Definition 18.6.14 (Day and Kelly [1970], Scott [1972]) *In a complete lattice A, an upward closed set U is called* **Scott open** *if it is inaccessible by joins of ideals, i.e.*

- *if I is an ideal contained in $\overline{U}$, then $\bigsqcup I$ is also in $\overline{U}$.*

Symmetrically, a downward closed set D is called **Scott closed** *if it is closed under joins of ideals, i.e.*

- *if I is an ideal contained in D, then $\bigsqcup I$ is also in D.*

The names used in the previous definition are justified by the following observation.

Proposition 18.6.15 *The Scott open sets on a complete lattice form a topology, called the* **Scott topology***.*

Proof. Obviously, both $\emptyset$ and A are Scott open. It is thus enough to verify closure under finite intersections and arbitrary unions.

Given two Scott open sets U_1 and U_2, $U_1 \cap U_2$ is obviously still upward closed. To show that it is inaccessible by joins of ideals, suppose I is an ideal such that $\bigsqcup I \in U_1 \cap U_2$. For $i = 1, 2$, $\bigsqcup I \in U_i$, and by inaccessibility of U_i, $I \not\subseteq \overline{U_i}$, i.e. there is $x_i \in I \cap U_i$. Now $x_1 \sqcup x_2 \in I$ because I is an ideal, and $x_1 \sqcup x_2 \in U_i$ because it is above x_i, and U_i is closed upward. Then $I \cap U_1 \cap U_2 \neq \emptyset$, i.e. $I \not\subseteq U_1 \cap U_2$.

Given a family $\{U_i\}_{i \in J}$ of Scott open sets, $\bigcup_{i \in J} U_i$ is obviously still upward closed.[10] To show it is inaccessible by joins of ideals, suppose I is an ideal contained in $\overline{\bigcup_{i \in J} U_i} = \bigcap_{i \in J} \overline{U_i}$. Then $I \subseteq \overline{U_i}$ for every $i \in J$, and by inaccessibility of U_i, $\bigsqcup I \in \overline{U_i}$, i.e. $\bigsqcup I \in \bigcap_{i \in J} \overline{U_i} = \overline{\bigcup_{i \in J} U_i}$. □

Exercises 18.6.16 a) *The complement of a principal ideal, i.e. a set of the form $\{x : x \not\sqsubseteq a\}$, is Scott open.* (Hint: a principal ideal is downward closed and closed under joins of ideals, i.e. it is Scott closed.)

b) *In a linear ordering, a non trivial subset is Scott open if and only if it is the complement of a principal ideal, i.e. a left-open upward interval.* (Hint: a Scott closed set must be closed under joins of arbitrary chains, and thus it must have a greatest element.)

[10] Notice that the union of filters is not necessarily a filter: this is one of the reasons why Scott open sets are defined in general for upward closed sets, and not only for filters.

c) *In a finite lattice, a subset is Scott open if and only if it is upward closed.* (Hint: ideals are closed under finite joins, and in a finite lattice they are closed under arbitrary joins.)

d) *In a power set, a subset is Scott open if and only if it is a union of principal filters generated by finite sets.* (Hint: one direction follows from the fact that every set is the join of the ideal of its finite subsets. For the other direction, suppose $\bigcup I \in U$, where I is an ideal, and U a union of principal filters generated by finite sets u. Then $\bigcup I$ is in some of these principal filters, i.e. there is a finite set $u \subseteq U$ such that $u \subseteq \bigcup I$. Then u must be contained in a finite subset of I, and hence be in I by the closure properties of ideals.)

Note that from part a) of the previous exercises it follows that *a Scott topology is* T_0, since if $x \not\sqsubseteq y$, then $\{z : z \not\sqsubseteq y\}$ is a Scott open set containing x but not y.

Moreover, *from a Scott topology we can recover the order of its underlying space*, as follows:

$$x \sqsubseteq y \iff (\forall U \text{ Scott open})(x \in U \Longrightarrow y \in U).$$

Indeed, if $x \sqsubseteq y$ and $x \in U$, then $y \in U$ by upward closure of U. And if $x \not\sqsubseteq y$, then $\{z : z \not\sqsubseteq y\}$ is a Scott open set containing x but not y.

It follows that *a Scott topology is not* T_2, unless the lattice has only one element. Indeed, if $x \neq y$ in a T_2 topology, then there are disjoint open sets separating x and y, and so $x \not\sqsubseteq y$ and $y \not\sqsubseteq x$.

Having defined a topology, the next step is to look at the notion of continuity induced by it. Recall that, given two topological spaces A and B, a function $f : A \to B$ is continuous if, for every open subset X of B, $f^{-1}(X)$ is an open subset of A. The next result relates this notion of continuity to one more familiar to us.

Proposition 18.6.17 (Scott [1972]) *The following are equivalent, for functions on complete lattices:*

1. *f is Scott continuous, i.e. continuous w.r.t. the Scott topology*
2. *f preserves l.u.b.'s of ideals, i.e. for every ideal I*

$$f(\bigsqcup I) = \bigsqcup\{f(i) : i \in I\} = \bigsqcup f(I).$$

Proof. To prove that 1 implies 2, we first show:

- a Scott continuous function is monotone

  By 18.6.16.a, $X = \{z : z \not\sqsubseteq f(y)\}$ is a Scott open subset of B, and so $f^{-1}(X)$ is a Scott open subset of A, hence closed upward. Suppose $f(x) \not\sqsubseteq f(y)$. Then $f(x) \in X$, i.e. $x \in f^{-1}(X)$. If $x \sqsubseteq y$, then $y \in f^{-1}(X)$ by upward closure of $f^{-1}(X)$, i.e. $f(y) \in X$, contradiction. Then $x \not\sqsubseteq y$.

Let now I be an ideal of A. By monotonicity of f,

$$\bigsqcup f(I) \sqsubseteq f(\bigsqcup I).$$

For the converse, consider the subset of B

$$X = \{z : z \not\sqsubseteq \bigsqcup f(I)\},$$

which is Scott open by 18.6.16.a. Then $f^{-1}(X)$ is a Scott open subset of A, hence inaccessible by joins of ideals. Notice that $I \subseteq \overline{f^{-1}(X)}$, because if $i \in I$, then $f(i) \sqsubseteq \bigsqcup f(I)$, i.e. $f(i) \notin X$, and $i \notin f^{-1}(X)$. Then $\bigsqcup I \in \overline{f^{-1}(X)}$, i.e. $f(\bigsqcup I) \notin X$, and so $f(\bigsqcup I) \sqsubseteq f(I)$.

To prove that 2 implies 1, we first show:

- *a function preserving l.u.b.'s of ideals is monotone*

  Indeed, $\{x : x \sqsubseteq y\}$ is an ideal with l.u.b. y, and thus

  $$f(y) = f(\bigsqcup\{x : x \sqsubseteq y\}) = \bigsqcup\{f(x) : x \sqsubseteq y\}.$$

  It follows that if $x \sqsubseteq y$, then $f(x) \sqsubseteq f(y)$.

Suppose now $f : A \to B$ preserves l.u.b.'s of ideals, and X is a Scott open subset of B. We prove that $f^{-1}(X)$ is a Scott open subset of A, so that f is continuous.

- *upward closure*

  If $x \sqsubseteq y$, then $f(x) \sqsubseteq f(y)$ by monotonicity of f. If $x \in f^{-1}(X)$, then $f(x) \in X$, and so $f(y) \in X$ by upward closure of X. Thus $y \in f^{-1}(X)$.

- *inaccessibility by joins of ideals*

  Let I be an ideal contained in $\overline{f^{-1}(X)}$. Then the downward closure J of $f(I)$ is an ideal contained in $\overline{X}$. Indeed, if x and y are in J, then by definition of J there are x_0 and y_0 in I such that $x \sqsubseteq f(x_0)$ and $y \sqsubseteq f(y_0)$. By monotonicity of f,

  $$x \sqcup y \sqsubseteq f(x_0) \sqcup f(y_0) \sqsubseteq f(x_0 \sqcup y_0),$$

  so that $x \sqcup y \in J$, and J is closed under $\sqcup$.

  Then, since f preserves l.u.b.'s of ideals and $f(I) \subseteq J$,

  $$f(\bigsqcup I) = \bigsqcup f(I) \sqsubseteq \bigsqcup J.$$

  But X is inaccessible by joins of ideals, so $\bigsqcup J \in \overline{X}$. Then $f(\bigsqcup I) \in \overline{X}$, and $\bigsqcup I \in \overline{f^{-1}(X)}$. □

The proof above shows, in particular, that *a Scott continuous function on complete lattices is monotone*.

In the proofs of 18.6.5 and 18.6.9 use was made of the way below relation $\ll$ to obtain Scott open filters. This is typical of continuous lattices, due to the next result.

Proposition 18.6.18 Scott Topology of Continuous Lattices (Scott [1972]) *On continuous lattices:*

1. *A subset X is Scott open if and only if*

$$X = \bigcup_{a \in X} \{x : a \ll x\}.$$

2. *A function f is Scott continuous if and only if, for every x,*

$$f(x) = \bigsqcup \{f(a) : a \ll x\}.$$

Proof. Part 1 is proved by the following:

- $\{x : a \ll x\}$ *is Scott open*

  It is upward closed because if $a \ll x$ and $x \sqsubseteq y$, then $a \ll y$.

  It is inaccessible by joins of ideals because if I is an ideal such that $a \ll \bigsqcup I$, then $a \in I$ by definition of $\ll$.

- *if X is Scott open, then $X = \bigcup_{a \in X} \{x : a \ll x\}$*

  The $\supseteq$ part follows from the fact that if $a \ll x$, then $a \sqsubseteq x$. Indeed, if $a \in X$ and $a \ll x$, then $x \in X$ because X is upward closed.

  To show the $\subseteq$ part, given $x \in X$, it is enough to find $a \in X$ such that $a \ll x$. Suppose no such a exists. Then $\{a : a \ll x\}$ is an ideal I (by definition of $\ll$) contained in $\overline{X}$ (by assumption) and such that $x = \bigsqcup I$ (by continuity). But $x \in X$, so $\bigsqcup I \in X$, and X is accessible by joins of ideals, contradiction.

To prove part 2, first suppose f is Scott continuous. Then it preserves l.u.b.'s of ideals by 18.6.17, and since $\{a : a \ll x\}$ is an ideal with l.u.b. x,

$$f(x) = f(\bigsqcup \{a : a \ll x\}) = \bigsqcup \{f(a) : a \ll x\}.$$

For the opposite direction, we first show:

- *f is monotone*

  Suppose $x \sqsubseteq y$. Then $\{a : a \ll x\} \subseteq \{a : a \ll y\}$, and

$$f(x) = \bigsqcup \{f(a) : a \ll x\} \sqsubseteq \bigsqcup \{f(a) : a \ll y\} = f(y).$$

If I is an ideal, then $\bigsqcup f(I) \sqsubseteq f(\bigsqcup I)$ follows from monotonicity. Conversely, by definition of $\ll$ we have that if $a \ll \bigsqcup I$, then $a \in I$. Hence

$$f(\bigsqcup I) = \bigsqcup\{f(a) : a \ll \bigsqcup I\} \sqsubseteq \{f(a) : a \in I\} = \bigsqcup f(I). \quad \Box$$

The main tool of the proofs of 18.6.5 and 18.6.9 can now be rephrased as saying that *if $x \not\sqsubseteq y$ in a continuous lattice, then there is a Scott open filter containing x but not y.*

Exercises 18.6.19 Scott topology of algebraic lattices. (Scott [1972]) Algebraic lattices were defined in 18.5.1.

a) *On algebraic lattices, a subset X is Scott open if and only if*

$$X = \bigcup_{a \in X} \{x : a \textit{ compact } \wedge\ a \sqsubseteq x\}.$$

(Hint: by 18.6.3.a, if a is compact, then $a \ll a$, and thus $\{x : a \sqsubseteq x\} = \{x : a \ll x\}$ is Scott open. Conversely, if X is Scott open and $x \in X$, suppose there is no compact element $a \in X$ such that $a \sqsubseteq x$. Then the downward closure of $\{a : a \text{ compact} \wedge a \sqsubseteq x\}$ is an ideal contained in $\overline{X}$ and with l.u.b. x, by 18.5.5.a and because the lattice is algebraic.)

b) *On algebraic lattices, a function f is Scott continuous if and only if*

$$f(x) = \bigsqcup\{f(a) : a \textit{ compact } \wedge\ a \sqsubseteq x\}.$$

(Hint: on algebraic lattices, $a \ll x$ if and only if there is a compact element k such that $a \sqsubseteq k \sqsubseteq x$.)

We have introduced the Scott topology of a complete lattice, and showed how the notion particularizes to continuous (and algebraic) lattices. It is also possible to go in the opposite direction, and generalize the notion to partial orderings that are not necessarily lattices.

Since the definition of an ideal requires only the existence of $\sqsubseteq$ and $\sqcup$, there is no trouble in extending the theory to uppersemilattices in which every ideal has a l.u.b.'s.

For partial orderings that are not necessarily uppersemilattices, we first have to rephrase the notion of ideal. This is achieved by the following trick.

Definition 18.6.20 *A subset D of a partial ordering is* **directed** *if, for every x and y in D, there is z in D such that $x, y \sqsubseteq z$. In other words, every finite subset of D is bounded in D.*

The idea is now to replace ideals by downward closed, directed sets. Actually, downward closure plays no role as far as l.u.b.'s are concerned, and thus we can restrict attention to the following notion, already introduced in 6.3.22.

Definition 18.6.21 *A partially ordered set is a* **directed complete partial ordering (d.c.p.o.)** *if every directed subset of it has a l.u.b. in it.*

We can now extend Definition 18.6.14 as follows.

Definition 18.6.22 *In a d.c.p.o., an upward closed set U is called* **Scott open** *if it is inaccessible by joins of directed sets, i.e.*

- *if D is a directed set contained in $\overline{U}$, then $\bigsqcup D$ is also in $\overline{U}$.*

To get an analogue of continuous lattices, some additional work is still needed. First we extend the definition 18.6.1 of $\ll$ as follows.

Definition 18.6.23 *An element a of a partial ordering is* **way below** *another element x ($\boldsymbol{a \ll x}$) if a is in the downward closure of any directed set D such that $x \sqsubseteq \bigsqcup D$.*

While the intersection of ideals is still an ideal, the intersection of directed sets is not necessarily directed, even if they are downward closed. The reason is that if D_1 and D_2 are downward closed directed sets, and $x \sqcap y \in D_1 \cap D_2$, then for $i = 1, 2$ there is $z_i \in D_i$ such that $x, y \sqsubseteq z_i$, but nothing ensures that there is $z \in D_1 \cap D_2$ such that $x, y \sqsubseteq z$.

To state the notion of continuity we thus have first to ensure that the appropriate sets are directed, so that their l.u.b.'s exist in a d.c.p.o.

Definition 18.6.24 (Markowsky [1976], [1981]) *A partial ordering is* **continuous** *if it is a d.c.p.o. and, for all x:*

1. *$\{a : a \ll x\}$ is directed*
2. *$x = \bigsqcup \{a : a \ll x\}$.*

Exercise 18.6.25 *A d.c.p.o. is continuous if and only if the operation $\bigsqcup$ on downward closed, directed sets has a left adjoint, i.e. there is f such that*

$$f(a) \subseteq D \iff a \sqsubseteq \bigsqcup D.$$

(Hint: if f exists, then $f(a) = \{x : x \ll a\}$ because it is contained in every downward closed, directed set D such that $a \sqsubseteq \bigsqcup D$. Thus $\{x : x \ll a\}$ is directed, and if $D = f(a)$, then a is the l.u.b. of D. Conversely, if the d.c.p.o. is continuous it is enough to let $f(a) = \{x : x \ll a\}$.)

At this point the reader can easily check that the theory developed in this subsection for complete lattices extends to d.c.p.o.'s. More precisely:

- The Scott open sets on a d.c.p.o. form a topology.

- A function on d.c.p.o's is Scott continuous if and only if it preserves l.u.b.'s of directed sets.
- On continuous p.o.'s, a subset X is Scott open if and only

$$X = \bigcup_{a \in X} \{x : a \ll x\},$$

and a function f is Scott continuous if and only if, for every x,

$$f(x) = \bigsqcup \{f(a) : a \ll x\}.$$

Chains of elements are special directed sets, and a further generalization from d.c.p.o.'s to c.c.p.o.'s produces the theory developed in Chapter 6, in particular the notion of continuity introduced in 6.3.7, whose connections with Scott topology were already discussed in 6.3.9.

Exercises 18.6.26 Topological Characterization of Completely Distributive Heyting Algebras. (Lawson [1979], Hofmann [1981])

a) *Every algebra of open sets of the Scott topology of a continuous lattice is a completely distributive Heyting algebra.* (Hint: by 18.6.13.b, it is enough to show that the algebra of open sets of the Scott topology of a continuous lattice is continuous and with enough co-points.)

b) *There are completely distributive Heyting algebras that are not isomorphic to Scott topologies of continuous lattices.* (Hint: suppose that the Boolean algebra with four elements is the algebra of open sets of a topological space X. In particular, $0 = \emptyset$ and $1 = X$. If the topology is the Scott topology of a lattice, from it we can recover the ordering. Then X can have at most three elements, two of which incomparable and the third their g.l.b.)

c) *Every completely distributive Heyting algebra is isomorphic to the algebra of open sets of the Scott topology of a continuous p.o.* (Hint: by 18.6.5 and 18.3.3, a completely distributive Heyting algebra A has enough points. By the proof of 18.3.4, it is then isomorphic to the algebra of open sets of the form $p(x) = \{a : a \text{ point } \wedge\ a \not\sqsupseteq x\}$. The crucial observation is that, by 18.6.16.a and the fact that $\mathrm{Pt}(A) \subseteq A$, each $p(x)$ is Scott open on the p.o. $(\mathrm{Pt}(A), \sqsupseteq)$ (notice the reverse ordering). It is then enough to check that the dual Stone topology of $(A, \sqsubseteq)$ coincides with the Scott topology of $(\mathrm{Pt}(A), \sqsupseteq)$, i.e. that $(\mathrm{Pt}(A), \sqsupseteq)$ is a continuous p.o. and every open set X of its Scott topology is a union of $p(x)$'s.)

d) *The completely distributive Heyting algebras are, up to isomorphism, exactly the algebras of open sets of Scott topologies of continuous p.o.'s.* (Hint: by part a) extended to continuous p.o.'s, and part c).)

18.7 Refinements of the Completeness Theorem

The Algebraic Completeness Theorem tells us that the class of all Heyting algebras is complete for the Intuitionistic Propositional Calculus, in the sense that if a formula is not provable, then there is a Heyting algebra in which it fails. Having introduced a whole spectrum of Heyting algebras, we now look for improvements.

Finite Heyting algebras

The next result shows that the class of finite Heyting algebras is enough for the Algebraic Completeness Theorem, as we already stated in 5.2.4.

Proposition 18.7.1 Finite Model Property (Jaskowski [1936], McKinsey and Tarski [1946]) *If $\Gamma \models_a \alpha$ fails, then there is a finite Heyting algebra $\mathcal{A}$ and an environment ρ on it such that all formulas of Γ are evaluated to 1 under it, but α is not.*

Proof. If $\Gamma \models_a \alpha$ fails, then by the Algebraic Completeness Theorem there is a Heyting algebra $\mathcal{B}$ and an environment η on it such that all all formulas of Γ are evaluated to 1 under it, but α is not. We have to find a finite Heyting algebra $\mathcal{B}$ with the same properties.

The first idea would be to consider all subformulas of Γ and α, and the Heyting subalgebra of $\mathcal{B}$ generated by their interpretations, i.e. by

$$\{[\![\beta]\!]^{\mathcal{B}}_{\eta} : \beta \text{ is a subformula of } \Gamma \cup \{\alpha\}\}.$$

The problem with this is that a Heyting algebra generated by a finite set of generators is not necessarily finite, since there is in general no way to collapse terms built with $\Rightarrow$.

To salvage at least part of the idea, we can certainly consider the sublattice $\mathcal{A}$ of $\mathcal{B}$ generated by the interpretations of subformulas of $\Gamma \cup \{\alpha\}$ (plus 0 and 1). $\mathcal{A}$ is finite because $\mathcal{B}$ is distributive, and the idempotency, associative, commutative and distributive laws allow us to collapse all but finitely many terms built with $\sqcap$ and $\sqcup$ from a finite number of elements (more precisely, from n elements we can at most obtain 2^n distinct subsets, and hence 2^n distinct meets, and then 2^{2^n} distinct joins of them).

$\mathcal{A}$ is a distributive lattice, being a sublattice of the distributive lattice $\mathcal{B}$. Then $\mathcal{A}$ is a finite distributive lattice, and hence a Heyting algebra, with its own right adjoint function $\Rightarrow_{\mathcal{A}}$.

Consider any environment ρ on $\mathcal{A}$ coinciding with η on the letters of $\Gamma \cup \{\alpha\}$. To show that $\mathcal{A}$ and ρ have the needed properties, i.e. that all formulas of Γ are evaluated to 1 under ρ, but α is not, it is enough to show that, for any subformula β of $\Gamma \cup \{\alpha\}$,

$$[\![\beta]\!]^{\mathcal{A}}_{\rho} = [\![\beta]\!]^{\mathcal{B}}_{\eta}.$$

This is almost immediate by induction on β, since ρ, $\sqcap_{\mathcal{A}}$, $\sqcup_{\mathcal{A}}$ (and $\sqsubseteq_{\mathcal{A}}$) coincide with η, $\sqcap_{\mathcal{B}}$, $\sqcup_{\mathcal{B}}$ (and $\sqsubseteq_{\mathcal{B}}$) when needed.

The proviso ‘almost’ refers to the fact that we still have to show that $\Rightarrow_{\mathcal{A}}$ and $\Rightarrow_{\mathcal{B}}$ also coincide when needed. This is the only non trivial point, since $\mathcal{A}$ is not necessarily closed under $\Rightarrow_{\mathcal{B}}$. We thus prove that, for elements $a, b \in \mathcal{A}$,

$$(a \Rightarrow_{\mathcal{A}} b) = (a \Rightarrow_{\mathcal{B}} b).$$

By definition of right adjoint:

1. For any $x \in \mathcal{A}$, $(x \sqcap a) \sqsubseteq b \iff x \sqsubseteq (a \Rightarrow_{\mathcal{A}} b)$.
2. For any $x \in \mathcal{B}$, $(x \sqcap a) \sqsubseteq b \iff x \sqsubseteq (a \Rightarrow_{\mathcal{B}} b)$.

We should have used $\sqcap_{\mathcal{A}}$ and $\sqsubseteq_{\mathcal{A}}$ in 1, and $\sqcap_{\mathcal{B}}$ and $\sqsubseteq_{\mathcal{B}}$ in 2, but they pairwise coincide on elements of $\mathcal{A}$ because the latter is a sublattice of $\mathcal{B}$, and thus no confusion arises.

Since $(a \Rightarrow_{\mathcal{A}} b) \in \mathcal{A}$, we can let $x = (a \Rightarrow_{\mathcal{A}} b)$ in 1. Then the right-hand-side automatically holds, and hence so does the left-hand-side, i.e. $[(a \Rightarrow_{\mathcal{A}} b) \sqcap a] \sqsubseteq b$. By 2, $(a \Rightarrow_{\mathcal{A}} b) \sqsubseteq (a \Rightarrow_{\mathcal{B}} b)$.

Since $(a \Rightarrow_{\mathcal{B}} b) \in \mathcal{B}$, we can let $x = (a \Rightarrow_{\mathcal{B}} b)$ in 2. Then the right-hand-side automatically holds, and hence so does the left-hand-side, i.e. $[(a \Rightarrow_{\mathcal{B}} b) \sqcap a] \sqsubseteq b$. By 1, $(a \Rightarrow_{\mathcal{B}} b) \sqsubseteq (a \Rightarrow_{\mathcal{A}} b)$. □

Although the whole class of finite Heyting algebras suffices for the Algebraic Completeness Theorem, no single finite Heyting algebra does. This stands in contrast with the case of the Classical Propositional Calculus (see 20.2.5).

Proposition 18.7.2 Failure of the Strong Algebraic Completeness Theorem (Gödel [1932]) *There is no single finite Heyting algebra $\mathcal{A}$ such that, if $\models_a \alpha$ fails, then there is an environment ρ on $\mathcal{A}$ such that α is not evaluated to* 1 *under it.*

Proof. We first consider the case of a Heyting algebra with two elements, i.e. the Boolean algebra $\{0, 1\}$. We could easily dispense with this case by just noticing that if α is any classically true but intuitionistically unprovable formula (e.g. Peirce’s Law), then $\models_a \alpha$ fails, but α is evaluated to 1 under any assignment on $\{0, 1\}$. The fact is that this proof does not generalize to algebras with more than two elements, and we thus consider an alternative one.

The idea is that, since $\{0, 1\}$ only has two elements, given any three distinct letters, at least two of them will have to agree under any assignment (obviously, not necessarily the same ones under different assignments). Let us thus consider the formula

$$(p \leftrightarrow q) \vee (p \leftrightarrow r) \vee (q \leftrightarrow r).$$

By definition of canonical interpretation, if ρ is an environment on $\{0, 1\}$, then $[\![\alpha]\!]_\rho = 1$. Indeed, at least two letters must get the same value under ρ, so at least one disjunct is evaluated to 1, and then so is the whole formula.

It only remains to note that α is not intuitionistically provable. Suppose it is. By the Disjunction Property, so would be one of its disjuncts, e.g. $p \leftrightarrow q$. But this is impossible, because such a disjunct is not even classically provable, not being a tautology (more precisely, being false when the two letters have different truth-values).

It is now clear how to generalize the proof to the case of Heyting algebras with n elements. It is enough to consider $n+1$ different letters $p_1, \ldots, p_{n+1}$, and the formula

$$\bigvee_{i \neq j} (p_i \leftrightarrow p_j).$$

As above, such a formula is not intuitionistically provable, but is evaluated to 1 under any environment on a Heyting algebra with n elements. □

Topological Heyting algebras

Although no single finite Heyting algebra suffices for the Algebraic Completeness Theorem, some infinite ones do. The first example is artificially constructed, but is the best possible in the spectrum of Heyting algebras introduced in the previous section.

Proposition 18.7.3 (Kripke [1963]) *There is a single complete Heyting algebra with enough strong co-points $\mathcal{A}$ such that, if $\models_a \alpha$ fails, then there is an environment ρ on $\mathcal{A}$ such that α is not evaluated to* 1 *under it.*

Proof. By 2.2.9 there is a Kripke model in which every formula α that is not intuitionistically provable is not forced. By 5.3.10, associated with such a Kripke model there is a topological Heyting algebra and an environment on it that produces a canonical interpretation coinciding with forcing. By 5.3.9 such a topology, being associated with a partial ordering, is closed under arbitrary intersections. By 18.2.4, such a topological Heyting algebra is thus complete and with enough strong co-points. □

The next example shows that the most common topological spaces are already complete, and provides nontrivial examples in the next classes of the spectrum of Heyting algebras.

Proposition 18.7.4 (Tarski [1938]) *Let $\mathcal{A}$ be the algebra of open sets of a metric space without isolated points. If $\Gamma \models_a \alpha$ fails, then there is an environment ρ on $\mathcal{A}$ such that all formulas of Γ are evaluated to* 1 *under it, but α is not.*

Proof.
□

Examples of metric spaces without isolated points are:

- the Baire space ω^ω, which is complete with enough points but not continuous (see 18.6.5)
- the euclidean space $I\!R$, which is continuous but not algebraic (see 18.6.4)
- the Cantor space 2^ω, which is algebraic without strong co-points (see 18.5.3).

æ

Part F

Classical Propositional Calculus

Chapter 19

Classical Propositional Calculus

The main reason to restrict attention to intuitionistic systems of propositional calculus is the nice correspondence with systems of typed λ-calculus, through the Curry-Howard isomorphism.

However, the intuitionistic (or, more generally, the constructive) approach to logic is not currently regarded as the standard one, and classical logic is most commonly used in mathematical practice. We thus take a detour to deal with it, that will also shed light on the relationships beween the two approaches.

19.1 Classical Implication

Having seen that any of the systems $\mathcal{N}$, $\mathcal{H}$ and $\mathcal{S}$ captures the notion of intuitionistic validity, we are left with a question: which modifications are needed to capture the notion of classical validity? This turns out to be quite easy to answer.

We start with a modification of the sequent system, from which the Classical Completeness Theorem is obtained naturally, and we then *discover* which addional axioms are needed to prove the equivalence with Natural Deduction and Hilbert systems.

The nocounterexample interpretation

The basic idea is to look at the notion of logical validity in the contrapositive, by means of the socalled **nocounterexample interpretation**: since $\Gamma \models \alpha$ means that α is true in all worlds in which all formulas in Γ are, *a proof of* $\Gamma \models \alpha$ *can be seen as the record of an unsuccessful attempt to describe a counterexample*, i.e.

a world in which this fails. Thus we start with the assertion that all formulas in Γ are true and α is false, and analyze the possible consequences of it.

The problem is that analyzing a truth assertion for a *single* formula may produce *two* truth assertions about its components: $\alpha \to \beta$ is true if α is false or β is true; but it is false if α is true *and* β is false. Inductively, an attempt to describe a counterexample may in general contain assertions about finitely many formulas being true and finitely many being false.

The general notion to deal with is thus the symmetric $\mathbf{\Gamma} \models \mathbf{\Delta}$, in which finite sets of formulas can appear both on the left and on the right, with the following intended interpretation: *there is no world in which all formulas of* Γ *are true and all formulas of* Δ *are false*. To get a Completeness Theorem we set up a formal system with rules that, when used backwards on $\Gamma \models \Delta$, allow us to replace a truth or falsity assertion about a given formula in Γ or Δ by truth or falsity assertions about its components. We can then systematically apply the rules, until all formulas on both sides are reduced to letters and cannot be further analyzed. We are thus left with the description of a number of possible worlds that would provide a counterexample to $\Gamma \models \Delta$. If any of such descriptions is consistent, i.e. no letter is required to be true and false at the same time, then we have the wanted counterexample. Otherwise, we have proved that no such counterexample exists. Inconsistent descriptions of worlds can thus be taken as axioms in this system, and a proof of $\Gamma \models \Delta$ consists of a tree starting from such axioms, proceeding by forward application of the rules, and ending with $\Gamma \models \Delta$.

The interesting point is that, when writing down the rules for $\Gamma \models \Delta$ in the way just described, we obtain rules similar to those of $\mathcal{S}$, the only difference being that we now consider sequents with possibly more than one formula on the right. In other words, we *discover* that this extension of sequents makes the associated system $\mathcal{SC}$ compatible with a semantical interpretation, radically different from the computational interpretation used in Section 1.3: $\Gamma \vdash_{\mathcal{S}} \Delta$ can now be taken to mean that there is no world in which all formulas of Γ are true and all formulas of Δ are false.

Definition 19.1.1 (Gentzen [1935]) *The relation $\vdash_{\mathcal{SC}}$ is inductively defined as follows:*

1. **Assumptions**. *Assumptions are sequents in which one formula appears on both sides of $\vdash_{\mathcal{SC}}$:*
$$\Gamma, \beta \vdash_{\mathcal{SC}} \beta, \Delta.$$

2. **$\to$-Introduction on the right**. *If β is deducible from Γ and α, then $\alpha \to \beta$ is deducible from Γ:*
$$\frac{\Gamma, \alpha \vdash_{\mathcal{SC}} \beta, \Delta}{\Gamma \vdash_{\mathcal{SC}} \alpha \to \beta, \Delta.}$$

3. **$\rightarrow$-Introduction on the left**. *If α is deducible from Γ and γ is deducible from Γ and β, then γ is deducible from Γ:*

$$\frac{\Gamma \vdash_{\mathcal{SC}} \alpha, \Delta \qquad \Gamma, \beta \vdash_{\mathcal{SC}} \Delta}{\Gamma, \alpha \rightarrow \beta \vdash_{\mathcal{SC}} \Delta.}$$

Notice that the rules just introduced are still *backward deterministic*, and thus the **Subformula Property** still holds: *in a proof of a sequent $\Gamma \vdash_{\mathcal{SC}} \Delta$, only subformulas of formulas in Γ or Δ can occur*.

The classical system with cut is defined by the natural modification of the Cut Rule.

Definition 19.1.2 Cut Rule. *The system $\mathcal{SC}$ + Cut is defined as the system $\mathcal{SC}$, with the additional rule:*

$$\frac{\Gamma \vdash_{\mathcal{SC}+\mathrm{Cut}} \gamma, \Delta \qquad \Gamma, \gamma \vdash_{\mathcal{SC}+\mathrm{Cut}} \Delta}{\Gamma \vdash_{\mathcal{SC}+\mathrm{Cut}} \Delta.}$$

The Cut Elimination procedure of Section 1.3 can be easily adapted to prove the following result.

Theorem 19.1.3 Cut Elimination (Gentzen [1935]) *For any Γ and β:*

$$\Gamma \vdash_{\mathcal{SC}+\mathrm{Cut}} \beta \;\Rightarrow\; \Gamma \vdash_{\mathcal{SC}} \beta.$$

A version of the semantical proof of Cut Elimination is also possible, using the Soundness and Completeness Theorem proved below, with a much simpler proof than in the intuitionistic case.

We now formally prove that the system $\mathcal{SC}$ really captures classical validity. This is not surprising, since the rules of the system were designed to make this result work (actually, the proof has already been sketched before 19.1.1). The surprise is rather that the simple change from one to finitely many formulas on the right makes the notion of a sequent compatible with a radically different interpretation, that accomodates classical reasoning.

Theorem 19.1.4 Classical Soundness and Completeness (Post [1921]) *For any Γ and Δ:*

$$\Gamma \vdash_{\mathcal{SC}} \Delta \;\Leftrightarrow\; \Gamma \models \Delta.$$

Proof. The left to right, soundness direction is done inductively on 19.1.1, and is like the proof of the Classical Soundness Theorem 2.1.3.

The right to left, completeness direction is done by contrapositive, by proving that if $\Gamma \vdash_{\mathcal{SC}} \Delta$, then there is a world $\mathcal{A}$ such that all formulas of Γ are true in $\mathcal{A}$, and all formulas of Δ are false in $\mathcal{A}$.

We first build any potential proof of $\Gamma \vdash_{\mathcal{SC}} \Delta$ by: starting from the latter as root; working upwards by using the rules of $\rightarrow$-introduction on the right or on the left, in any order; stopping at one node when only propositional letters remain on each side of the sequent. The procedure obviously halts, because at every step the complexity of one formula, either on the right or on the left, decreases. If every leaf is an assumption, i.e. if a same letter occurs on both sides of $\vdash_{\mathcal{SC}}$, then we reached a proof because each such sequent is an assumption.

Otherwise, consider any leaf whose associated sequent is not an assumption. Let $\mathcal{A}$ be the set of letters appearing on the left: such a world agrees with the sequent, in the sense that it makes true the formulas (in this case: letters) on the left by definition, and false the ones on the right by the classical definition of truth 2.1.1, since no letter appears on both sides. We now show, by induction on 19.1.1, that it also agrees with every sequent on the path from it to the root. In particular, it agrees with the root, and hence it makes all formulas of Γ true and all of Δ false, so that $\Gamma \not\models \Delta$.

For $\rightarrow$-introduction on the right, suppose $\mathcal{A}$ agrees with $\Gamma, \alpha \vdash_{\mathcal{SC}} \beta, \Delta$. Then it makes all formulas of Γ and α true, and all formulas of Δ and β false. In particular, it makes $\alpha \rightarrow \beta$ false, and thus it agrees with $\Gamma \vdash_{\mathcal{SC}} \alpha \rightarrow \beta, \Delta$.

For $\rightarrow$-introduction on the left, we have two cases. First, suppose that $\mathcal{A}$ agrees with $\Gamma, \beta \vdash_{\mathcal{SC}} \Delta$. Then it makes β true, and hence $\alpha \rightarrow \beta$ true. Thus $\mathcal{A}$ agrees with $\Gamma, \alpha \rightarrow \beta \vdash_{\mathcal{SC}} \Delta$.

Second, suppose that $\mathcal{A}$ agrees with $\Gamma \vdash_{\mathcal{SC}} \alpha, \Delta$. Then it makes α false, and hence $\alpha \rightarrow \beta$ true. Thus $\mathcal{A}$ agrees with $\Gamma, \alpha \rightarrow \beta \vdash_{\mathcal{SC}} \Delta$.[1] □

Classical tableaux

The sequent system for the Intuitionistic Implicational Calculus is more efficient than the tableaux method: basically, sound restrictions (producing only finite proofs) of the two methods are much easier to state for the former (see 3.1.2) than for the latter (see Nerode [1990]).

In the classical case tableaux are just reformulations of sequent proofs, and thus the two methods are in this case practically the same. A slight advantage of the former is that we do not have to drag hypotheses along, since they are recorded on the branches of the tableaux, and thus proofs can be written in a less cumbersome way.

Definition 19.1.5 *A* **classical tableau** *is a tree with nodes consisting of signed forcing assertions of the form $T\alpha$ or $F\alpha$, and consistent with the following formation rules:*

[1] Notice that this is the case that fails in the intuitionistic system $\mathcal{S}$, where Δ is empty in the hypothesis but not in the conclusion: if we only know that $\mathcal{A}$ agrees with $\Gamma \vdash_{\mathcal{S}} \alpha$, then we do not know that it also agrees with $\Gamma, \alpha \rightarrow \beta \vdash_{\mathcal{S}} \gamma$, because we do not know anything about γ.

1. *If a node $T\alpha \to \beta$ is on the tree, then we can split any branch going through it by adding $F\alpha$ in one direction and $T\beta$ in the other. Graphically,*

$$\begin{array}{c} T\alpha \to \beta \\ \hline\hline F\alpha \quad T\beta, \end{array}$$

where the double line shows that the bottom nodes do not have to immediately follow the top one.

2. *If a node $F\alpha \to \beta$ is on the tree, then we can extend any branch going through it by adding $T\alpha$ and $F\beta$. Graphically,*

$$\begin{array}{c} F\alpha \to \beta \\ \hline\hline T\alpha \\ \hline F\beta. \end{array}$$

Notice the asymmetric treatment: in the first case we split branches, in the second case we linearly extend them. Classical tableaux are simply a reformulation of $\mathcal{SC}$, since the two rules above correspond, repsectively, to →-introduction on the left and on the right.

Definition 19.1.6 Δ *is* **provable by classical tableaux** *from* Γ *(written* $\boldsymbol{\Gamma} \vdash_{\mathcal{TC}} \boldsymbol{\Delta}$*) if there is a classical tableau starting from $T\gamma$ for all $\gamma \in \Gamma$ and $F\delta$ for all $\delta \in \Delta$, such that all its branch are contradictory, in the sense that on every branch there is a pair of nodes of the form $T\beta$ and $F\beta$ (the same β for any given branch, although possibly different β's for different branches).*

The next definition captures the idea of systematic search.

Definition 19.1.7 *A* **complete systematic tableau** *is a tableau in which the rules have been used exaustively, in the sense that:*

1. *For any node $T\alpha \to \beta$ on the tree and any branch going through it, there is a node on the branch that splits into two nodes $F\alpha$ and $T\beta$.*

2. *For any node $F\alpha \to \beta$ on the tree and any branch going through it, there is a node on the branch followed by two nodes $T\alpha$ and $F\beta$.*

In an actual construction of a complete systematic tableau, it is enough to apply the rules relative to any given node to any branch going through it, and to mark off the relative node. Since the rules replace nodes relative to a given formula by nodes relative to subformulas of it, after finitely many steps the only **unmarked nodes** will be those relative to atomic formulas. In particular, *a complete systematic tableau is finite*.

Obviously, whenever in the construction of a complete systematic tableau we hit a contradiction along a branch, we can seal that branch off and stop developing it, since every extension of it will remain contradictory.

As an example of the method, we prove Pierce's Law:

$$\begin{array}{c}
F[(p \to q) \to p] \to p \\ \hline
T(p \to q) \to p \\ \hline
Fp \\ \hline
\begin{array}{cc}
\begin{array}{c} Fp \to q \\ \hline Tp \\ \hline Fq \end{array} & \begin{array}{c} Tp. \\ \\ \\ \end{array}
\end{array}
\end{array}$$

Notice that this tableau is a finitary version of the intuitionistic tableau for Pierce's Law considered on p. 44: in the latter the three final nodes on the left branch were repeated infinitely often, each time w.r.t. to a new extension.

The proof of the next result shows that, as already announced, $\vdash_{\mathcal{TC}}$ is simply a reformulation of $\vdash_{\mathcal{SC}}$. Basically, *a tableau is a sequent tree upside down, and the (unmarked nodes of the) branches of the former correspond to the (formulas of the sequents on the) leaves of the latter.*

Proposition 19.1.8 *There are canonical translations of classical sequent proofs to classical tableaux proofs, and conversely.*

Proof. Given a classical tableau, we proceed by induction on its construction and build a sequent tree as follows. If the tableau starts from $T\gamma$ for $\gamma \in \Gamma$ and $F\delta$ for $\delta \in \Delta$, we let $\Gamma \vdash_{\mathcal{SC}} \Delta$ be the root node of the sequent tree. At any stage, if rule

$$\begin{array}{c}
T\alpha \to \beta \\ \hline\hline
F\alpha \quad T\beta
\end{array}$$

is applied at a given node, then we extend the tree above the corresponding node according to the rule

$$\frac{\Gamma_1 \vdash_{\mathcal{SC}} \alpha, \Delta_1 \qquad \Gamma_1, \beta \vdash_{\mathcal{SC}} \Delta_1}{\Gamma_1, \alpha \to \beta \vdash_{\mathcal{SC}} \Delta_1.}$$

Similarly, if rule

$$\begin{array}{c}
F\alpha \to \beta \\ \hline\hline
T\alpha \\ \hline
F\beta
\end{array}$$

is applied at a given node, we extend the tree above the corresponding node according to the rule

$$\frac{\Gamma_1, \alpha \vdash_{\mathcal{SC}} \beta, \Delta_1}{\Gamma_1 \vdash_{\mathcal{SC}} \alpha \to \beta, \Delta_1.}$$

By induction, each node $\Gamma_1 \vdash_{\mathcal{SC}} \Delta_1$ on the tree thus built records the unmarked nodes $T\alpha$ and $F\beta$ (for $\alpha \in \Gamma_1$ and $\beta \in \Delta_1$) of the part of the corresponding branch of the tableau considered so far. If the tableau is a proof, then any branch is contradictory and there is a formula β such that $T\beta$ and $F\beta$ are on it. The corresponding sequent is then an axiom of $\mathcal{SC}$, since the same formula appears on both sides of $\vdash_{\mathcal{SC}}$, and thus the sequent tree is a sequent proof.

Conversely, given a sequent tree we proceed in a symmetrical way by induction on the definition of $\vdash_{\mathcal{SC}}$ and build a tableau such that, at any given stage, a sequent $\Gamma_1 \vdash_{\mathcal{SC}} \Delta_1$ records the unmarked nodes $T\alpha$ and $F\beta$ (for $\alpha \in \Gamma_1$ and $\beta \in \Delta_1$) of the corresponding branch of the tableau built thus far. If the tree is a sequent proof, then every leaf is an axiom and there is a formula β appearing on both sides $\vdash_{\mathcal{SC}}$. The corresponding branch of the tableau is thus contradictory, and the tableau is a proof. □

Corollary 19.1.9 Equivalence of Classical Sequents and Tableaux (Beth [1955], Hintikka [1955]) *For any* Γ *and* Δ*,*

$$\Gamma \vdash_{\mathcal{TC}} \Delta \;\Leftrightarrow\; \Gamma \vdash_{\mathcal{SC}} \Delta.$$

It follows from the Classical Soundness and Completeness Theorem that

$$\Gamma \vdash_{\mathcal{TC}} \Delta \;\Leftrightarrow\; \Gamma \models \Delta.$$

A direct proof of this result can be copied down from the proof of 19.1.4, using the translation described in 19.1.8. In particular, from any noncontradictory branch of any complete systematic tableau starting from $T\gamma$ for all $\gamma \in \Gamma$, and $F\delta$ for all $\delta \in \Delta$, we can read off a classical world $\mathcal{A}$ in which all $\gamma \in \Gamma$ are true and all $\delta \in \Delta$ are false. Precisely, $\mathcal{A}$ is the set of all letters p such that Tp is on the branch.

19.2 Classical Propositional Calculus

Sequents and tableaux

difference with the intuitionistic case: sets of formulas in the consequence.

Functional completeness and definability

Disjunctive normal form.

$\neg$ and $\wedge$ or $\rightarrow$.

$\rightarrow$ and $\wedge$ (and $\vee$?) are not adequate (by induction, formulas built from them and p are either equivalent to p or always true. Every formula is satisfiable (by the assignment that makes all letters true).

19.3 Complexity

Decidability

The decision procedure for $\mathcal{SC}$ is a consequence of the Subformula Property, as in the intuitionistic case 3.1.2. The classical case is even easier, in two respects. On the one hand, *a sequent can have only finitely many cut-free proofs*. On the other hand, *every possible analysis of →-introduction produces a proof*.

Proposition 19.3.1 *The relation* $\vdash_{\mathcal{SC}}$ *is decidable.*

Proof. By the proof of 19.1.4, to decide whether $\Gamma \vdash_{\mathcal{SC}} \Delta$ it is enough to build any potential proof by: starting from the latter as root; working upwards by using the rules of →-introduction on the right or on the left, in any order; stopping at one node when only propositional letters remain on each side of the sequent. The procedure halts by the Subformula Property and, by the proof of 19.1.4, $\Gamma \vdash_{\mathcal{SC}} \Delta$ holds if and only if the potential proof is an actual proof. □

The decision procedure for $\mathcal{TC}$ is similar, and slightly more efficient. The reason is that tableaux are easier to write down, since they do not require dragging hypotheses along.

Proposition 19.3.2 *The relation* $\vdash_{\mathcal{TC}}$ *is decidable.*

Proof. To decide whether $\Gamma \vdash_{\mathcal{SC}} \Delta$ it is enough to build any complete systematic tableau starting from $T\gamma$ for $\gamma \in \Gamma$, and $F\delta$ for $\delta \in \Delta$. Then, by the proofs of 19.1.4 and 19.1.9, $\Gamma \vdash_{\mathcal{SC}} \Delta$ holds if and only if the tableau is contradictory. □

The decision procedure for $\models$ is a consequence of the fact, easily proved by induction, that classical worlds agreeing on the letters occurring in a given formula also agree on the formula itself.

Proposition 19.3.3 *The relation* $\models$ *is decidable.*

Proof. By definition, $\Gamma \models \alpha$ holds if and only if α is true in every classical world $\mathcal{A}$ in which all formulas of Γ are true. This apparently requires the consideration of all possible classical worlds $\mathcal{A}$, and there are infinitely many of them. But the definition of truth is inductive, and thus only the finitely many letters $p_1, \ldots, p_n$ occurring in α or in some formulas of Γ matter. Then it is enough to consider every possible combination of membership values for $p_1, \ldots, p_n$ in $\mathcal{A}$ (there are 2^n possible ones), and check that whenever one such combination makes all formulas of Γ true, it also makes α true. □

A standard and simple way of carrying out the previous decision procedure is by building a **truth-table** with $m+n+1$ columns (respectively corresponding to the letters $p_1, \ldots, p_n$ of α, the formulas $\gamma_1, \ldots, \gamma_m$ of Γ, and α itself) and 2^n rows (corresponding to the restrictions to $p_1, \ldots, p_n$ of all possible worlds). Letters T and F are taken to mean 'true' and 'false', respectively, and are placed under the letters $p_1, \ldots, p_n$ in all of their 2^n possible combinations. Then a T or an F is eventually placed under (the main implication of) γ_i's and α, by inductively using the rule stated by the following truth-table (defining implication):

β	γ	$\beta \to \gamma$
T	T	T
T	F	F
F	T	T
F	F	T

Then $\Gamma \models \alpha$ holds if there is a T under α in any row in which there is a T under every γ_i. In particular, $\models \alpha$ holds if the column under the main implication sign of α consists only of T's.

For example, the following truth-table proves the validity of an instance of Axiom 2 of $\mathcal{H}$:

p	q	p	$\to$	$(q \to p)$
T	T		T	T
T	F		T	T
F	T		T	F
F	F		T	T

Among the decision procedures proposed above, the most efficient is the one based on tableaux. In particular, the truth-table method is always exponential in the number of letters occurring in a given formula.

Complexity

As alredy noted, classical decidability is apparently less complicated than intuitionistic decidability. This is made precise by the next result, that measures the complexity of the latter.

!!! State the following for the full propositional calculus, and then deduce it for the implicational calculus !!!

Theorem 19.3.4 (Cook [1972]) *The complexity of $\models$ is co-NP-complete.*

Proof. □

æ

Chapter 20

Boolean Algebras

In this chapter we look at Boolean algebras and their connections with the Classical Propositional Calculus. Sections 1–3 follow the blueprint of Chapter 5 for Heyting algebras. Section 4 and 5 look at Boolean algebras from the points of view of Heyting algebras and cartesian closed categories.

Since the theory of Boolean algebras is highly developed, we will only touch here on aspects directly related to our main interest. For a general view of the subject the reader can turn to Monk and Bonnet [1989].

20.1 Boolean Algebras

Complements

Since $\alpha \vee \neg\alpha$ is a tautology and $\alpha \wedge \neg\alpha$ is a contradiction, the behaviour of the canonical interpretation of classical negation is captured by the following notion.

Definition 20.1.1 *An element a of a Heyting algebra is complemented if there is an element b such that*

$$a \sqcup b = 1 \qquad \textit{and} \qquad a \sqcap b = 0.$$

b is called the **complement** *of a, and is indicated by $\overline{a}$.*

The following proposition shows that it makes sense to talk of *the* complement of an element in Heyting algebras (which, we recall, are distributive by 17.1.24).

Proposition 20.1.2 *In a distributive lattice complements are unique, when they exist.*

Proof. Suppose both b_1 and b_2 are complements of a. Then

$$\begin{array}{rcl} b_1 & = & b_1 \sqcap 1 \\ & = & b_1 \sqcap (a \sqcup b_2) \\ & = & (b_1 \sqcap a) \sqcup (b_1 \sqcap b_2) \\ & = & 0 \sqcup (b_1 \sqcap b_2) \\ & = & b_1 \sqcap b_2 \end{array}$$

because b_2 is a complement, by distributivity, and because b_1 is a complement. Then $b_1 \sqsubseteq b_2$. Symmetrically, $b_2 \sqsubseteq b_1$, and thus $b_1 = b_2$. □

The next result shows that $\overline{}$ behaves in a Heyting algebra as $\neg$ does in classic logic.

Proposition 20.1.3 *The following laws hold in any Heyting algebra, for complemented elements:*

1. **Excluded Middle.** $a \sqcup \overline{a} = 1$
2. **Double Negation.** $\overline{\overline{a}} = a$.
3. **De Morgan.** $\overline{a \sqcup b} = \overline{a} \sqcap \overline{b}$ *and* $\overline{a \sqcap b} = \overline{a} \sqcup \overline{b}$
4. **Contrapositive.** $a \sqsubseteq b$ *if and only if* $\overline{b} \sqsubseteq \overline{a}$

Proof. 1 is an obvious consequence of the definition of $\overline{}$.

2 follows from the fact that if $a \sqcup b = 1$ and $a \sqcap b = 0$, then b is the complement of a and a is the complement of b, i.e. $b = \overline{a}$ and $a = \overline{b} = \overline{\overline{a}}$.

To prove 3, notice that

$$\begin{array}{rcl} (a \sqcup b) \sqcap (\overline{a} \sqcap \overline{b}) & = & [(a \sqcup b) \sqcap \overline{a}] \sqcap \overline{b} \\ & = & [(a \sqcap \overline{a}) \sqcup (b \sqcap \overline{a})] \sqcap \overline{b} \\ & = & [0 \sqcup (b \sqcap \overline{a})] \sqcap \overline{b} \\ & = & b \sqcap \overline{a} \sqcap \overline{b} \\ & = & 0 \end{array}$$

by associativity, distributivity, and the facts that $a \sqcap \overline{a} = 0$ and $b \sqcap \overline{b} = 0$. Symmetrically,

$$(a \sqcup b) \sqcup (\overline{a} \sqcap \overline{b}) = 1.$$

Thus $\overline{a} \sqcap \overline{b}$ is the complement of $a \sqcup b$. Similarly, $\overline{a} \sqcup \overline{b}$ is the complement of $a \sqcap b$.

To prove 4, it is enough to prove the left to right direction, since the right to left follows from it (applied to $\overline{b} \sqsubseteq \overline{a}$), by the Double Negation Law. Suppose that $a \sqsubseteq b$. Then $a \sqcup b = b$, and by the De Morgan Laws

$$\overline{a} \sqcap \overline{b} = \overline{a \sqcup b} = \overline{b},$$

i.e. $\overline{b} \sqsubseteq \overline{a}$. □

Boolean algebras

Definition 20.1.4 *A* **Boolean algebra** *is a complemented Heyting algebra, i.e. a Heyting algebra in which every element is complemented.*

The next result provides an alternative and self-contained approach to Boolean algebras, which avoids any reference to Heyting algebras and the adjointness condition defining $\Rightarrow$.

Proposition 20.1.5 *The Boolean algebras are exactly the complemented distributive lattices.*

Proof. Any Boolean algebra is a distributive lattice, because so is any Heyting algebra.

Conversely, it is enough to show that any complemented distributive lattice can be turned into a Heyting algebra, by letting

$$(a \Rightarrow b) = (\overline{a} \sqcup b).$$

By definition of adjointness, we need to show

$$x \sqcap a \sqsubseteq b \iff x \sqsubseteq \overline{a} \sqcup b.$$

If $x \sqcap a \sqsubseteq b$, then

$$x \sqsubseteq \overline{a} \sqcup x = (\overline{a} \sqcup x) \sqcap (\overline{a} \sqcup a) = \overline{a} \sqcup (x \sqcap a) \sqsubseteq \overline{a} \sqcup b$$

because $\overline{a} \sqcup a = 1$, and by distributivity.

If $x \sqsubseteq \overline{a} \sqcup b$, then

$$x \sqcap a \sqsubseteq (\overline{a} \sqcup b) \sqcap a = (\overline{a} \sqcap a) \sqcup (b \sqcap a) = b \sqcap a \sqsubseteq b$$

by distributivity, and because $\overline{a} \sqcap a = 0$. □

Notice that the two conditions of being complemented and distributive serve dual purposes in a lattice: the first ensures the existence of complements, the second their uniqueness.

Exercises 20.1.6 Boolean rings (Stone [1935]). A **Boolean ring** is a ring with identity $\langle A, +, \cdot\,, 0, 1\rangle$ such that $a^2 = a$ for every a, i.e. every element is idempotent.

a) *Any Boolean algebra is a Boolean ring.* (Hint: let

$$a + b = (a \sqcap \overline{b}) \sqcup (\overline{a} \sqcap b) \qquad \text{and} \qquad a \cdot b = a \sqcap b.)$$

b) *Any Boolean ring is a Boolean algebra.* (Hint: let

$$a \sqcup b = a + b + a \cdot b, \qquad a \sqcap b = a \cdot b \qquad \text{and} \qquad \overline{a} = 1 + a.)$$

c) *The correspondence between Boolean algebras and Boolean rings given by parts a) and b) is a bijection.*

d) *Boolean rings are commutative.* (Hint:

$$a + b = (a + b)^2 = a^2 + a \cdot b + b \cdot a + b^2 = a + a \cdot b + b \cdot a + b,$$

so $a \cdot b + b \cdot a = 0$ and $a \cdot b = -b \cdot a$. By letting $a = b$ we get $a^2 = -a^2$, i.e. $a = -a$ for any a. Then $a \cdot b = b \cdot a$.)

Examples

We follow the blueprint of Section 5.3, where we gave examples of Heyting algebras.

On the negative side, we have the following result.

Proposition 20.1.7 *The only linear orderings that are Boolean algebras are the trivial ones, i.e.* $\{0\}$ *and* $\{0, 1\}$.

Proof. If $x \sqsubseteq y$, then $x \sqcup y = y$ and $x \sqcap y = x$. Thus the only complemented elements of a linear ordering are 0 and 1, when they exist. □

In particular, any non trivial linear ordering with 0 and 1 is an example of *a Heyting algebra which is not a Boolean algebra.*

On the positive side, we provide two classes of examples trivially satisfying the conditions of being complemented distributive lattices, but crucial for the representation theorems of Section 20.3

The first class is purely set-theoretical.

Proposition 20.1.8 *Any field of sets, i.e. a field whose elements are sets and whose sum and product are the set-theoretical intersection and union, is a Boolean algebra.*

In particular, any power set ordered under set-theoretical inclusion is a Boolean algebra.

The second class is topological, and requires the following definitions (in addition to the ones of topology and open set given on p. 85):

- a *closed set* in a topology is a set whose set-theoretical complement is open
- a *clopen set* is a set which is both closed and open or, equivalently, a set which is open together with its set-theoretical complement.

Proposition 20.1.9 *The clopen sets of a topology ordered under set-theoretical inclusion form a Boolean algebra.*

20.2 Soundness and Completeness Theorem

Definition 20.2.1 *A formula α is a* **classical algebraic consequence** *of $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ (written $\boldsymbol{\Gamma \models_{ac} \alpha}$) if for every Boolean algebra $\mathcal{A}$ and every environment ρ on it,*

$$(\llbracket\gamma_1\rrbracket_\rho^{\mathcal{A}} \sqcap_{\mathcal{A}} \cdots \sqcap_{\mathcal{A}} \llbracket\gamma_n\rrbracket_\rho^{\mathcal{A}}) \sqsubseteq_{\mathcal{A}} \llbracket\alpha\rrbracket_\rho^{\mathcal{A}}.$$

In the limit case of Γ empty, we get the notion of classical algebraic validity: α is **classically algebraically valid** (written $\models_{ac} \alpha$) if α evaluates to $1_{\mathcal{A}}$ in every Boolean algebra, under every environment.

Lindenbaum algebras

Theorem 20.2.2 Algebraic Soundness and Completeness (Boole [1854], Post [1921], Tarski [1935a]) *For any Γ and α,*

$$\Gamma \vdash_{\mathcal{NC}} \alpha \iff \Gamma \models_{ac} \alpha.$$

Proof. As in 5.2.3 and 17.1.26, we consider the structure

$$\mathcal{A}_\Gamma = \langle A_\Gamma, \sqsubseteq, =, \sqcap, \sqcup, \Rightarrow, ^-, 0, 1\rangle$$

in which:

1. A_Γ is the set of equivalence classes

$$\llbracket\beta\rrbracket = \{\gamma : \Gamma \vdash_{\mathcal{NC}} \beta \leftrightarrow \gamma\}$$

2. $\sqsubseteq$ is induced by $\vdash_{\mathcal{NC}}$ relatively to Γ, i.e.

$$\llbracket\beta\rrbracket \sqsubseteq \llbracket\gamma\rrbracket \iff \Gamma, \beta \vdash_{\mathcal{NC}} \gamma$$

3. $=$ is induced by provable equivalence relatively to Γ, i.e.

$$\llbracket\beta\rrbracket = \llbracket\gamma\rrbracket \iff \Gamma \vdash_{\mathcal{NC}} (\beta \leftrightarrow \gamma)$$

4. $\sqcap$ is induced by $\wedge$, i.e.
$$[\![\beta]\!] \sqcap [\![\gamma]\!] = [\![\beta \wedge \gamma]\!]$$

5. $\Rightarrow$ is induced by $\rightarrow$, i.e.
$$[\![\beta]\!] \Rightarrow [\![\gamma]\!] = [\![\beta \rightarrow \gamma]\!]$$

6. $^-$ is induced by $\neg$, i.e.
$$\overline{[\![\beta]\!]} = [\![\neg\beta]\!]$$

7. 0 is the equivalence class of the formulas refutable from Γ, i.e.
$$0 = \{\beta : \Gamma \vdash_{\mathcal{NC}} \neg\beta\}$$

8. 1 is the equivalence class of the formulas provable from Γ, i.e.
$$1 = \{\beta : \Gamma \vdash_{\mathcal{NC}} \beta\}.$$

We already know from 17.1.26 that $\mathcal{A}_\Gamma$ is a Heyting algebra, and to prove that it is a Boolean algebra it is enough to show that it is complemented. Since $^-$ is induced by $\neg$, this follows from

$$\vdash_{\mathcal{NC}} \alpha \vee \neg\alpha \qquad \text{and} \qquad \vdash_{\mathcal{NC}} \neg(\alpha \wedge \neg\alpha),$$

which translate into

$$[\![\alpha]\!] \sqcup \overline{[\![\alpha]\!]} = 1 \qquad \text{and} \qquad [\![\alpha]\!] \sqcap \overline{[\![\alpha]\!]} = 0.$$

This proves
$$\Gamma \models_{ac} \alpha \implies \Gamma \vdash_{\mathcal{NC}} \alpha.$$

Conversely, to prove
$$\Gamma \vdash_{\mathcal{NC}} \alpha \implies \Gamma \models_{ac} \alpha,$$

we have to show that if all formulas in Γ are true in a Boolean algebra, then so is α. But since
$$\Gamma \vdash_{\mathcal{NC}} \alpha \iff \Gamma, \text{Excluded Middle} \vdash_{\mathcal{N}} \alpha,$$

it is enough to notice that a Boolean algebra in which all formulas in Γ are true, is a Heyting algebra in which the Excluded Middle and all formulas of Γ are true (see 20.4.15 for details). By the Intuitionistic Algebraic Soundness Theorem 17.1.26, it then follows that α is true, too. □

The Algebraic Completeness Theorem provides us with a canonical Boolean algebra $\mathcal{A}_\emptyset$, consisting of the equivalence classes of formulas under the equivalence relation induced by classical provable equivalence.

The Algebraic Soundness Theorem shows that any function from the propositional letters to a Boolean algebra $\mathcal{A}$, i.e. any environment on $\mathcal{A}$, can be extended to a homomorphism of Boolean algebras from $\mathcal{A}_\emptyset$ to $\mathcal{A}$, i.e. to the canonical interpretation associated with the environment. This property is concisely expressed by saying that $\mathcal{A}_\emptyset$ is the *free Boolean algebra on countably many generators*. More precisely, the generators are the equivalence classes of propositional letters, which are countably many because distinct letters cannot be provably equivalent.

Exercise 20.2.3 *$\mathcal{A}_\emptyset$ is isomorphic to the Boolean algebra of the clopen sets of the Cantor space.* (Rasiowa and Sikorski [1963]) (Hint: recall that the Cantor space is the set $\mathcal{P}(\omega)$ of all subsets of ω, with the topology generated by the basic open sets $\{X : X \supseteq u\}$, with u finite. Since all free Boolean algebras on the same number of generators are isomorphic, it is enough to show that the clopen sets of the Cantor space are a free Boolean algebra on countably many generators. The generators are the basic open sets $\{X : x \in X\}$ with $x \in \omega$, i.e. the ones corresponding to singletons. Since each Boolean algebra is isomorphic to a field of sets by 20.3.2, it is enough to note that any function from the generators to a field of sets can be extended to a homomorphism of Boolean algebras by the natural set-theoretical operations.)

The two-element Boolean algebra

A version of the Finite Model Property 18.7.1 is easier to obtain for Boolean algebras than it was for Heyting algebras.

Exercise 20.2.4 Finite Model Property. *If $\Gamma \models_{ac} \alpha$ fails, then there is a finite Boolean algebra $\mathcal{A}$ and an environment ρ on it such that all formulas of Γ are evaluated to 1 under it, but α is not.* (Hint: as in the first part of the proof of 18.7.1, since the Laws of Idempotency, Associativity, Commutativity, Distributivity and the De Morgan's Law imply that a finitely generated Boolean algebra is finite.)

The reason to confine the Finite Model Property to the exercises is that a much stronger result actually holds here. More precisely, while for Heyting algebras the Finite Model Property could not be improved by finding a *single* finite model that worked in all cases, for Boolean algebras not only this is possible: the Boolean algebra with only *two* elements is enough. This provides a strong form of algebraic completeness.

Theorem 20.2.5 Strong Algebraic Completeness (Rasiowa and Sikorski [1950], Łoš [1951]) *If $\Gamma \models_{ac} \alpha$ fails, then there is an environment ρ on the Boolean algebra $\{0, 1\}$ such that all formulas of Γ are evaluated to 1 under it, but α is not.*

Proof. If $\Gamma \models_{ac} \alpha$ fails, by the Algebraic Completeness Theorem there is a Boolean algebra $\mathcal{A}$ and an environment ρ on it such that all formulas of Γ are evaluated to 1 under it, but α is not. Given a maximal filter F on $\mathcal{A}$ not containing $[\![\alpha]\!]_\rho$, consider the quotient Boolean algebra $\mathcal{A}_{/F}$, and the environment $\rho_{/F}$ induced by ρ on it. Since F is maximal, $\mathcal{A}_{/F} = \{0, 1\}$, and it has the needed properties. □

The next result shows that the Strong Algebraic Completeness Theorem is just a different formulation of the implication

$$\Gamma \models \alpha \implies \Gamma \models_{ac} \alpha,$$

and that the truth-table method is just a special case of an algebraic interpretation on the Boolean algebra $\{0, 1\}$.

Proposition 20.2.6 Canonical Algebraic Interpretation Induced by Classical Worlds. *Given a classical world $\mathcal{A}$, i.e. a set of propositional letters, the environment*

$$\rho(p) = \begin{cases} 1 & \textit{if } p \in \mathcal{A} \\ 0 & \textit{otherwise} \end{cases}$$

on the Boolean algebra $\{0, 1\}$ produces a canonical interpretation that coincides with truth in $\mathcal{A}$, in the sense that for every formula α

$$[\![\alpha]\!]_\rho = \begin{cases} 1 & \textit{if } \mathcal{A} \models \alpha \\ 0 & \textit{otherwise.} \end{cases}$$

Proof. We prove that

$$[\![\alpha]\!]_\rho = 1 \iff \mathcal{A} \models \alpha$$

by induction on α.

If $\alpha = p$, then this holds by definition of ρ.

If $\alpha = \neg\beta$, then

$$\begin{aligned} [\![\neg\beta]\!]_\rho = 1 &\iff \overline{[\![\beta]\!]_\rho} = 1 \\ &\iff [\![\beta]\!]_\rho = 0 \\ &\iff \mathcal{A} \not\models \beta \\ &\iff \mathcal{A} \models \neg\beta \end{aligned}$$

by definition of $[\![\ \ ]\!]$, the fact that 0 is the complement of 1, induction hypothesis, and definition of $\models$.

If $\alpha = \beta \wedge \gamma$, then

$$\begin{aligned} [\![\beta \wedge \gamma]\!]_\rho = 1 &\iff [\![\beta]\!]_\rho \sqcap [\![\gamma]\!]_\rho = 1 \\ &\iff [\![\beta]\!]_\rho = 1 \ \wedge\ [\![\gamma]\!]_\rho = 1 \\ &\iff \mathcal{A} \models \beta \ \wedge\ \mathcal{A} \models \gamma \\ &\iff \mathcal{A} \models \beta \wedge \gamma \end{aligned}$$

by definition of $[\![\ \ ]\!]$, the fact that 1 is the greatest element, induction hypothesis, and definition of $\models$.

The remaining cases are similar. □

Generic environments ⋆

20.2.5 and 20.2.6 say that the notion of classical validity is a special case of algebraic validity, since

$$\models_{ac} \alpha \iff (\forall \text{ Boolean algebra } \mathcal{A})(\forall \rho \text{ on } \mathcal{A})([\![\alpha]\!]_\rho^{\mathcal{A}} = 1)$$

and

$$\models \alpha \iff (\forall \rho \text{ on the fixed Boolean algebra } \{0,1\})([\![\alpha]\!]_\rho^{\mathcal{A}} = 1).$$

So, as usual, the Algebraic Soundness Theorem provides a *wider* class of models than the ones considered in the original notion of validity.

We now want to go in the opposite direction, and restrict the notion of classical validity by considering a *smaller* class of environments on $\{0,1\}$ that still provide completeness. This is obtained through Kripke forcing and the double negation interpretation, which we will prove in 21.1.3, as follows. Given any Kripke model

$$\mathcal{A} = \langle P_{\mathcal{A}}, \sqsubseteq_{\mathcal{A}}, \{\mathcal{A}_\sigma\}_{\sigma \in P_{\mathcal{A}}} \rangle,$$

any branch $\sigma_0 \sqsubseteq \sigma_1 \sqsubseteq \cdots$ determines the unique environment:

$$\rho(p) = \begin{cases} 1 & \text{if } p \in \bigcup \mathcal{A}_{\sigma_n} \\ 0 & \text{otherwise.} \end{cases}$$

Alternatively, we can think of $\bigcup \mathcal{A}_{\sigma_n}$ as providing a classical world that is a limit of intuitionistic ones, i.e. the set of propositional letters that eventually appear in the partial worlds associated to the given branch.

Definition 20.2.7 *A* **generic branch** $\sigma_0 \sqsubseteq \sigma_1 \sqsubseteq \cdots$ *of a Kripke model is a branch such that for every propositional formula α, there is an n such that $\sigma_n \Vdash \alpha$ or $\sigma_n \Vdash \neg\alpha$.*

ρ is a **generic environment** *if it arises from a generic branch of a Kripke model.*

Alternatively, one could define $\bigcup \mathcal{A}_{\sigma_n}$ to be a **generic (classical) world**.

Proposition 20.2.8 *For every Kripke model $\mathcal{A}$ and every σ on it, there is a generic branch extending σ.*

Proof. Let $\{\alpha_n\}_{n\in\omega}$ be an enumeration of the propositional formulas. Define $\sigma_0 = \sigma$, and at stage $n+1$ let σ_n be given. There are two possible cases:

- If $(\exists\sigma \sqsupseteq \sigma_n)(\sigma \Vdash \alpha_n)$, choose one such σ and let $\sigma_{n+1} = \sigma$, so that $\sigma_{n+1} \Vdash \alpha_n$.
- Otherwise, let $\sigma_{n+1} = \sigma_n$. Then $\sigma_{n+1} \Vdash \neg\alpha_n$ by definition of forcing for negation, since

$$\sigma_n \Vdash \neg\alpha_n \iff (\forall\sigma \sqsupseteq \sigma_n)(\sigma \not\Vdash \alpha_n). \quad \Box$$

In particular, generic environments exist.

Definition 20.2.9 *A formula α is a* **generic consequence** *of Γ (written $\boldsymbol{\Gamma \models_g \alpha}$) if, for every generic environment ρ on the Boolean algebra $\{0,1\}$, α is evaluated to 1 whenever all formulas of Γ are.*

Alternatively, one could define $\Gamma \models_g \alpha$ by saying that α is true in every generic world $\bigcup \mathcal{A}_{\sigma_n}$ in which every $\gamma \in \Gamma$ is true.

The connection between (intuitionistic) forcing and (classical) truth is given by the next result.

Proposition 20.2.10 *If $\{\sigma_n\}_{n\in\omega}$ is a generic branch of a Kripke model, and ρ is the generic environment associated to it, then for any formula α*

$$[\![\alpha]\!]_\rho = 1 \iff (\exists n)(\sigma_n \Vdash \alpha).$$

Proof. By induction on α. Genericity will take care of $\neg$, and $\wedge$ and $\vee$ are trivial because forcing for them is defined as classical truth. So the only non trivial case will be $\rightarrow$, because of the nonstandard definition of forcing for it.

If $\alpha = p$, then

$$\begin{aligned} [\![p]\!]_\rho = 1 &\iff \rho(p) = 1 \\ &\iff p \in \textstyle\bigcup \mathcal{A}_{\sigma_n} \\ &\iff (\exists n)(p \in \mathcal{A}_{\sigma_n}) \\ &\iff (\exists n)(\sigma_n \Vdash p) \end{aligned}$$

by definition of $[\![\ ]\!]$, of canonical interpretation, of union and of forcing for atomic formulas.

If $\alpha = \neg\beta$, then

$$\begin{aligned} [\![\neg\beta]\!]_\rho = 1 &\iff [\![\beta]\!]_\rho = 0 \\ &\iff (\forall n)(\sigma_n \not\Vdash \beta) \\ &\iff (\exists n)(\sigma_n \Vdash \neg\beta) \end{aligned}$$

by definition of $[\![\ ]\!]$, induction hypothesis and genericity.

If $\alpha = \beta \wedge \gamma$, then

$$\begin{aligned} [\![\beta \wedge \gamma]\!]_\rho = 1 &\iff [\![\beta]\!]_\rho = 1 \text{ and } [\![\gamma]\!]_\rho = 1 \\ &\iff (\exists n)(\sigma_n \Vdash \beta) \text{ and } (\exists m)(\sigma_m \Vdash \gamma) \\ &\iff (\exists p)(\sigma_p \Vdash \beta \text{ and } \sigma_p \Vdash \gamma) \\ &\iff (\exists p)(\sigma_p \Vdash \beta \wedge \gamma) \end{aligned}$$

by definition of $\llbracket\ \rrbracket$, induction hypothesis, monotonicity (take $p = \max\{m, n\}$) and definition of forcing.

If $\alpha = \beta \vee \gamma$, then

$$\begin{aligned}
\llbracket\beta \vee \gamma\rrbracket_\rho = 1 &\iff \llbracket\beta\rrbracket_\rho = 1 \text{ or } \llbracket\gamma\rrbracket_\rho = 1\\
&\iff (\exists n)(\sigma_n \Vdash \beta) \text{ or } (\exists n)(\sigma_n \Vdash \gamma)\\
&\iff (\exists n)(\sigma_n \Vdash \beta \text{ or } \sigma_n \Vdash \gamma)\\
&\iff (\exists n)(\sigma_n \Vdash \beta \vee \gamma)
\end{aligned}$$

by definition of $\llbracket\ \rrbracket$, induction hypothesis, properties of $\exists$ and definition of forcing.

If $\alpha = \beta \to \gamma$, then

$$\begin{aligned}
\llbracket\beta \to \gamma\rrbracket_\rho = 1 &\implies \llbracket\beta\rrbracket_\rho = 0 \text{ or } \llbracket\gamma\rrbracket_\rho = 1\\
&\implies (\forall n)(\sigma_n \nVdash \beta) \text{ or } (\exists n)(\sigma_n \Vdash \gamma)\\
&\implies (\exists n)(\sigma_n \Vdash \neg\beta) \text{ or } (\exists n)(\sigma_n \Vdash \gamma)\\
&\implies (\exists n)(\sigma_n \Vdash \neg\beta \text{ or } \sigma_n \Vdash \gamma)\\
&\implies (\exists n)[(\forall \tau \sqsupseteq \sigma_n)(\tau \nVdash \beta) \text{ or } (\forall \tau \sqsupseteq \sigma_n)(\tau \Vdash \gamma)]\\
&\implies (\exists n)(\forall \tau \sqsupseteq \sigma_n)(\tau \nVdash \beta \text{ or } \tau \Vdash \gamma)\\
&\implies (\exists n)(\forall \tau \sqsupseteq \sigma_n)(\tau \Vdash \beta \ \Rightarrow\ \tau \Vdash \gamma)\\
&\implies (\exists n)(\sigma_n \Vdash \beta \to \gamma)
\end{aligned}$$

by definition of $\llbracket\ \rrbracket$, induction hypothesis, genericity, properties of $\exists$, definition of forcing and monotonicity, properties of $\forall$, definition of $\Rightarrow$ and definition of forcing. And

$$\begin{aligned}
\llbracket\beta \to \gamma\rrbracket_\rho = 0 &\implies \llbracket\beta\rrbracket_\rho = 1 \text{ and } \llbracket\gamma\rrbracket_\rho = 0\\
&\implies (\exists n)(\sigma_n \Vdash \beta) \text{ and } (\forall n)(\sigma_n \nVdash \gamma)\\
&\implies (\exists n)(\sigma_n \Vdash \beta) \text{ and } (\exists m)(\sigma_m \Vdash \neg\gamma)\\
&\implies (\exists p)(\sigma_p \Vdash \beta \text{ and } \sigma_p \Vdash \neg\gamma)\\
&\implies (\exists p)[(\forall \tau \sqsupseteq \sigma_p)(\tau \Vdash \beta) \text{ and } (\forall \tau \sqsupseteq \sigma_p)(\tau \nVdash \gamma)]\\
&\implies (\exists p)(\forall \tau \sqsupseteq \sigma_p)(\tau \Vdash \beta \text{ and } \tau \nVdash \gamma)\\
&\implies (\exists p)(\forall \tau \sqsupseteq \sigma_p)(\tau \nVdash \beta \to \gamma)\\
&\implies (\forall n)(\sigma_n \nVdash \beta \to \gamma)
\end{aligned}$$

by definition of $\llbracket\ \rrbracket$, induction hypothesis, genericity, monotonicity (take $p = \max\{m, n\}$), monotonicity and definition of forcing, properties of $\forall$, definition of forcing and monotonicity. □

Theorem 20.2.11 Generic Completeness. *If $\Gamma \models \alpha$ fails, then there is a generic environment ρ on the Boolean algebra $\{0, 1\}$ such that all formulas of Γ are evaluated to 1 under it, but α is not.*

Proof. If $\Gamma \models \alpha$ fails, then so does $\neg\neg\Gamma \vdash_{\mathcal{N}} \neg\neg\alpha$ by the Double Negation interpretation 21.1.3. So there is a Kripke model on which all $\neg\neg\gamma$ with $\gamma \in \Gamma$ are forced, but $\neg\neg\alpha$ is not. By definition of forcing a double negation, this means that

$$(\forall\sigma)(\exists\tau \sqsupseteq \sigma)(\tau \Vdash \alpha)$$

fails, i.e.

$$(\exists\sigma)(\forall\tau \sqsupseteq \sigma)(\tau \nVdash \alpha).$$

Choose such a σ, and build a generic branch extending it as in 20.2.8. On such a branch:

- α is not forced by consistency of forcing, since $\sigma \Vdash \neg\alpha$ by definition, and thus $(\forall n)(\sigma_n \Vdash \neg\alpha)$, i.e. $(\forall n)(\sigma_n \nVdash \alpha)$;
- all $\gamma \in \Gamma$ are forced, since $\neg\neg\gamma$ is forced in the model, i.e.

  $$(\forall\sigma)(\exists\tau \sqsupseteq \sigma)(\tau \Vdash \gamma),$$

  and the construction of the generic branch will then force γ when it comes to it.

If ρ is the generic environment relative to that generic branch, then $[\![\gamma]\!]_\rho = 1$ for all $\gamma \in \Gamma$, but $[\![\alpha]\!]_\rho \neq 1$ (by the Truth Lemma), so $\Gamma \not\models \alpha$. □

Soundness and Completeness Theorems

We can now state the fundamental result about presentations of the Classical Propositional Calculus.

Theorem 20.2.12 The Magnificent Seven of Classical Propositional Calculus. *The following are equivalent, for any Γ and α:*

1. $\Gamma \vdash_{\mathcal{NC}} \alpha$ *(natural deduction)*
2. $\Gamma \vdash_{\mathcal{HC}} \alpha$ *(Hilbert system)*
3. $\Gamma \vdash_{\mathcal{SC}} \alpha$ *(sequent system)*
4. $\Gamma \vdash_{\mathcal{TC}} \alpha$ *(classical tableaux)*
5. $\Gamma \models_{ac} \alpha$ *(Boolean algebras)*
6. $\Gamma \models \alpha$ *(classical worlds)*
7. $\Gamma \models_g \alpha$ *(generic worlds).*

Proof. By 21.2.6, 20.2.2, 20.2.5 and 20.2.11. □

20.3 Representation Theorems ⋆

We now ask how far the examples of Boolean algebras produced in Section 20.1 are typical. Some of the results of this section are paradigms for the ones of Section 5.4 for Heyting algebras, and actually special cases of them, while others are specific for Boolean algebras.

Lindenbaum algebras

Theorem 20.3.1 First Representation for Boolean Algebras (Tarski [1935a]) *Any Boolean algebra is isomorphic to a Lindenbaum algebra for the Classical Propositional Calculus.*

Proof. As in 5.4.1, by working with Boolean algebras and $\vdash_{\mathcal{NC}}$ in place of Heyting algebras and $\vdash_{\mathcal{N}}$. □

Power sets

Theorem 20.3.2 Second Representation for Boolean Algebras (Stone [1936]) *Any Boolean algebra is isomorphic to a field of sets.*

Proof. The proofs of 5.4.2 and 17.1.28 already prove the result, since the function defined there has the needed properties. The only added piece of information is that f preserves complements too, and this is immediate because complements are characterized by the conditions

$$x \sqcup \overline{x} = 1 \qquad \text{and} \qquad x \sqcap \overline{x} = 0,$$

and $\sqcap$, $\sqcup$, 0, 1 are all preserved. □

Corollary 20.3.3 *Any Boolean algebra is isomorphic to a subalgebra of a power set.*

Proof. A field of sets is a subalgebra of the power set consisting of all subsets of its unit element. □

The formulation of the corollary just proved raises the question of which Boolean algebras are isomorphic not only to a *subalgebra*, but to a *full* power set. The answer is given in terms of the following analogue of 18.3.1.

Definition 20.3.4 (Schröder [1891]) *An element $a \neq 0$ of a Boolean algebra is an* **atom** *if*

$$x \sqsubseteq a \implies x = 0 \vee x = a,$$

i.e. if there is no other element between it and 0.

A Boolean algebra is **atomic** *if for every element $x \neq 0$ there is an atom a such that $a \sqsubseteq x$.*

The set of atoms of a Boolean algebra A is indicated by $\mathbf{At}(\boldsymbol{A})$.

Exercises 20.3.5 (Tarski [1935]) a) *An element $a \neq 0$ is an atom if and only if it is $\sqcup$-irreducible, i.e.*

$$a = (x \sqcup y) \implies (a = x) \vee (a = y).$$

(Hint: suppose a is $\sqcup$-irreducible and $0 \sqsubseteq x \sqsubseteq a$. Then

$$a = a \sqcap 1 = a \sqcap (x \sqcup \overline{x}) = (a \sqcap x) \sqcup (a \sqcap \overline{x}) = x \sqcup (a \sqcap \overline{x})$$

by distributivity, so $a = x$ or $a = a \sqcap \overline{x}$. In the latter case $x \sqsubseteq a \sqsubseteq \overline{x}$, i.e. $x = 0$.)

b) *An element $a \neq 0$ is an atom if and only if it is a co-point, i.e.*

$$a \sqsubseteq (x \sqcup y) \implies (a \sqsubseteq x) \vee (a \sqsubseteq y).$$

(Hint: by distributivity, as in 18.3.2.c.)

c) *An element $a \neq 0$ is an atom if and only if $\overline{a}$ is a point.* (Hint: the conditions

$$a \sqsubseteq x \sqcup y \implies a \sqsubseteq x \ \vee \ a \sqsubseteq y$$

and

$$\overline{a} \sqsupseteq \overline{x} \sqcap \overline{y} \implies \overline{a} \sqsupseteq \overline{x} \ \vee \ \overline{a} \sqsupseteq \overline{y}$$

are equivalent.)

d) *A Boolean algebra is atomic if and only if it has enough points as a Heyting algebra.* (Hint: first notice that a Boolean algebra A is atomic if and only if, for every x and y in A, if $x \neq y$, then there is an atom below one of x and y but not below the other. For example, if $x \not\sqsubseteq y$, then $x \sqcap \overline{y} \neq 0$, and an atom $a \sqsubseteq x \sqcap \overline{y}$ cannot be below y, otherwise $a \sqsubseteq y \sqcap \overline{y} = 0$.

Then notice that, for any atom a,

$$a \sqsubseteq x \iff \overline{a} \not\sqsupseteq x.$$

The left to right implication holds for any element $a \neq 0$. For the right to left implication, from $a \sqsubseteq 1 = x \sqcup \overline{x}$ we have

$$a = a \sqcap (x \sqcup \overline{x}) = (a \sqcap x) \sqcup (a \sqcap \overline{x}),$$

and by part a) either $a = a \sqcap x$ or $a = a \sqcap \overline{x}$, i.e. $a \sqsubseteq x$ or $a \sqsubseteq \overline{x}$. Then, by taking complements, $a \sqsubseteq x$ or $\overline{a} \sqsupseteq x$.)

e) *A Boolean algebra is atomic if and only if every element is the l.u.b. of the atoms below it.* (Hint: given $x \neq 0$, if x is not the l.u.b. of the atoms below it, then there is $y \sqsubset x$ above all such atoms. If A is atomic, by part d) there is an atom below one of x and y but not the other, contradiction.)

f) *Not every atomic Boolean algebra is complete, and not every complete Boolean algebra is atomic.* (Hint: consider the set of finite or cofinite elements of $\mathcal{P}(\omega)$, and the quotient of $\mathcal{P}(\omega)$ w.r.t. the ideal of finite sets or, equivalently, the filter of cofinite sets.)

The next result shows that for Boolean algebras the various notions of Heyting algebras studied in Chapter 17 all coincide among each other, as well as with the newly introduced notion of complete atomic Boolean algebra.

Proposition 20.3.6 (Tarski [1935], Lindenbaum) *The following are equivalent for a Boolean algebra A:*

1. *A is complete and atomic*
2. *A is complete with enough co-points*
3. *A is complete with enough strong co-points*
4. *A is algebraic*
5. *A is completely distributive*
6. *A is continuous*
7. *A is complete with enough points.*

Proof. Conditions 1, 2 and 7 are equivalent by 20.3.5.

Conditions 1 and 3 are equivalent in a similar way. In one direction, an atom is obviously $\bigsqcup$-irreducible, and hence a strong co-point by 18.2.2.c (notice that $\sqcap\bigsqcup$-distributivity holds in every Heyting algebra by 5.3.6, and hence in every Boolean algebra). In the opposite direction, a strong co-point is a co-point, and hence an atom.

Conditions 3 and 7 are equivalent by the first part of the proof. The remaining conditions 4, 5 and 6 are intermediate between conditions 3 and 7, and hence equivalent to them, by 18.5.3, 18.6.4, 18.6.5, 18.6.12.e and 18.6.12.f. $\square$

The previous result implies that the next one is a version of 18.3.4, but it is instructive to give a direct proof. Notice that for (complete) Heyting algebras we cannot use atoms in place of points, since in general the elements are not distinguished by the atoms below them (for example, a linear ordering with an atom is atomic, but all elements different from 0 have the same atoms below them). For Boolean algebras, instead, we can.

Theorem 20.3.7 (Tarski [1935], Lindenbaum) *A Boolean algebra is isomorphic to a power set if and only if it is atomic and complete.*

Proof. To show that the conditions are necessary, note that a power set is closed under arbitrary unions and intersections, its atoms are the singletons $\{x\}$, and every nonempty set (i.e. any element $\neq 0$) contains a singleton.

For sufficiency, let A be any Boolean algebra, and consider the function cp[1] from A to $\mathcal{P}(\mathrm{At}(A))$ defined as follows:

$$cp(x) = \{a : a \text{ atom } \wedge\ a \sqsubseteq x\}.$$

Then cp is automatically a homomorphism of Boolean algebras, for the following reasons:

- $cp(0) = \emptyset$
  By definition, for any atom a we have $a \neq 0$, and so $a \not\sqsubseteq 0$.

- $cp(1) =$ *the set of atoms of* A
  By definition, for any element a we have $a \sqsubseteq 1$.

- $cp(x \sqcap y) = cp(x) \cap cp(y)$
  For any element a,

  $$a \sqsubseteq (x \sqcap y) \iff (a \sqsubseteq x) \wedge (a \sqsubseteq y)$$

  by definition of $\sqcap$.

- $cp(x \sqcup y) = cp(x) \cup cp(y)$
  For any atom a,

  $$a \sqsubseteq (x \sqcup y) \iff (a \sqsubseteq x) \vee (a \sqsubseteq y).$$

  Indeed, the right to left implication holds by definition of $\sqcup$, for any element a. For the left to right implication, let a be an atom and $a \sqsubseteq x \sqcup y$. Then

  $$a = a \sqcap (x \sqcup y) = (a \sqcap x) \sqcup (a \sqcap y)$$

  by distributivity, and thus $a = x_1 \sqcup y_1$ for some $x_1 \sqsubseteq x$ and $y_1 \sqsubseteq y$. Since an atom is $\sqcup$-irreducible, $a = x_1$ or $a = y_1$, and hence $a \sqsubseteq x$ or $a \sqsubseteq y$.

- $cp(\overline{x}) = \overline{cp(x)}$
  This follows from the fact that complements are characterized by the conditions

  $$x \sqcup \overline{x} = 1 \qquad \text{and} \qquad x \sqcap \overline{x} = 0,$$

  and $\sqcap$, $\sqcup$, 0, 1 are all preserved.

  In particular, this shows that

  $$a \sqsubseteq \overline{x} \iff a \not\sqsubseteq x,$$

  for any atom a.

[1] We use the letters *cp* as a reminder that an atom is a co-point, see 20.3.5.b.

- *if A is atomic, then cp is one-one*
  If $x \neq y$, then $x \not\sqsubseteq y$ or $y \not\sqsubseteq x$. Suppose e.g. that $x \not\sqsubseteq y$. Then $x \sqcap \overline{y} \neq 0$. Otherwise $x \sqcap \overline{y} = 0$, and by distributivity

  $$\begin{aligned} y &= 0 \sqcup y \\ &= (x \sqcap \overline{y}) \sqcup y \\ &= (x \sqcup y) \sqcap (\overline{y} \sqcup y) \\ &= (x \sqcup y) \sqcap 1 \\ &= x \sqcup y, \end{aligned}$$

  contradicting $x \sqsubseteq y$.

  Since $x \sqcap \overline{y} \neq 0$ and A is atomic, there is an atom a such that $a \sqsubseteq x \sqcap \overline{y}$, and hence $a \sqsubseteq x$ and $a \sqsubseteq \overline{y}$. Then, as noticed above, $a \not\sqsubseteq y$, and so a is an atom below x but not below y, i.e. $cp(x) \neq cp(y)$.

- *if A is complete, then cp is onto*
  Given a set X of atoms, $\bigsqcup X$ exists because A is complete. We want to show that $cp(\bigsqcup X) = X$.

  The inclusion $X \subseteq cp(\bigsqcup X)$ is trivial: if $a \in X$, then obviously $a \sqsubseteq \bigsqcup X$, by definition of $\bigsqcup$.

  We prove the inclusion $cp(\bigsqcup X) \subseteq X$ by contrapositive: if $a \notin X$, we want to show $a \sqcap x = 0$, so that $a \not\sqsubseteq x$ (otherwise $0 = a \sqcap x = a \neq 0$).

  If $b \in X$, then $a \not\sqsubseteq b$ because both a and b are atoms, and $a \neq b$ since $a \notin X$ and $b \in X$. Then, as noticed above, $a \sqsubseteq \overline{b}$, and $a \sqcap b = 0$ (since $\overline{b} \sqcap b = 0$). So

  $$0 = \bigsqcup_{b \in X} (a \sqcap b) = a \sqcap (\bigsqcup_{b \in X} b) = a \sqcap x$$

  by complete distributivity. □

Finite Boolean algebras

The previous results have a number of trivial but nice consequences.

Theorem 20.3.8 Algebraic Characterization of Finite Boolean Algebras. *The finite Boolean algebras are, up to isomorphism, the finite fields of sets.*

Proof. By 20.1.8 and 20.3.2. □

Recall from **??** that, instead, the finite Heyting algebras are, up to isomorphism, the finite rings of set.

Theorem 20.3.9 Set-Theoretical Characterization of Finite Boolean Algebras. *The finite Boolean algebras are, up to isomorphism, the power sets of finite sets.*

Proof. By 20.1.8 and 20.3.7, since a finite Boolean algebra is obviously atomic and complete. □

Corollary 20.3.10 Classification Theorem for Finite Boolean Algebras. *Two finite Boolean algebras are isomorphic if and only if they have the same number of elements.*

Proof. The condition is obviously necessary. For sufficiency, by the previous characterization a finite Boolean algebra is the power set of a finite set, and hence it has 2^n elements and n atoms, for some n. Thus two finite Boolean algebras with the same number of elements must have the same number of atoms, and any bijection of their atoms induces an isomorphism between the two algebras. □

The Boolean Prime Ideal Theorem

Notice that in the proof of 20.3.7 we used the following properties of atoms:

- if $a \sqsubseteq x \sqcup y$, then $a \sqsubseteq x$ or $a \sqsubseteq y$
- $a \sqsubseteq x$ or $a \sqsubseteq \overline{x}$.

They respectively say that the principal filter generated by a is a *prime filter* and an *ultrafilter*. In a Boolean algebra the two conditions are equivalent, as the next exercises show.

Exercises 20.3.11 Ultrafilters. (Cartan [1937]) In a Boolean algebra a nontrivial filter F is called an **ultrafilter** if the following condition holds, for any x in A:

$$x \in F \quad \text{or} \quad \overline{x} \in F.$$

a) *A nontrivial filter is an ultrafilter if and only if it is prime.* (Hint: since $x \sqcup \overline{x} = 1 \in F$, by primality $x \in F$ or $\overline{x} \in F$, and so a prime filter is an ultrafilter. Conversely, if F is an ultrafilter and $x \notin F$ and $y \notin F$, then $\overline{x} \in F$ and $\overline{y} \in F$. So $\overline{x} \sqcap \overline{y} = \overline{x \sqcup y} \in F$ by closure under $\sqcap$. Thus $x \sqcup y \notin F$, otherwise F would contain both an element and its complement, hence their g.l.b. 0, and it would be trivial.)

b) *A non trivial filter is an ultrafilter if and only if it is maximal.* (Hint: let F be a maximal filter and $x \notin F$. The filter generated by $F \cup \{x\}$ is trivial by maximality, and so 0 is in it. By distributivity and definition of generated filter, there is $a \in F$ such that $a \sqcap x = 0$. Then $a \sqsubseteq \overline{x}$, and $\overline{x} \in F$ by upward closure. Conversely, let F be an ultrafilter, and suppose $x \notin F$. Then $\overline{x} \in F$, so both x and $\overline{x}$ are in the filter generated by $F \cup \{x\}$, which is trivial because $x \sqcap \overline{x} = 0$. Then F is maximal.)

c) *A subset X of a Boolean algebra A is an ultrafilter if and only if the characteristic function $c_X : A \to \{0,1\}$ is a homomorphism of Boolean algebras.* (Hint: the proof of 20.3.7 shows necessity. For sufficiency, e.g. if $c_X(\overline{x}) = \overline{c_X(x)}$, then $c_X(\overline{x}) = 0$ if and only if $c_X(x) = 1$, i.e. $x \in X$ if and only if $\overline{x} \notin X$. Similarly, $x \sqcup y \in X$ if and only if $x \in X$ or $y \in X$.)

If we replace the function

$$cp(x) = \{a : a \text{ atom } \wedge\ a \sqsubseteq x\}$$

by the function

$$u(x) = \{\mathcal{U} : \mathcal{U} \text{ is an ultrafilter containing } x\}$$

in the proof of 20.3.7, we still get a homomorphism of Boolean algebras. Moreover, the proof of the condition

$$\text{if } A \text{ is atomic, then } cp \text{ is one-one}$$

shows that

$$\text{if } A \text{ has enough ultrafilters, then } u \text{ is one-one.}$$

By the duality between (ultra)filters and (prime) ideals, a Boolean algebra has enough ultrafilters if and only if it has enough prime ideals. That any Boolean algebra has enough prime filters follows from the so-called **Boolean Prime Ideal Theorem** saying that if $x \not\sqsubseteq y$ on a Boolean algebra, then there is a prime ideal containing y but not x.[2] By verifying it, with a proof dual to that of 17.1.28, we get back the proof of 20.3.2, thus showing that the latter is a generalization of the one-one embedding part of 20.3.7 to arbitrary (not necessarily atomic) Boolean algebras. More precisely, the role of atoms is taken by ultrafilters. And the condition that there are sufficiently many atoms, i.e. that the Boolean algebra is atomic, becomes the condition that there are sufficiently many ultrafilters, which is always satisfied by the Boolean Prime Ideal Theorem.

Exercises 20.3.12 a) *The Boolean Prime Ideal is equivalent to the existence of an ultrafilter not containing any given $x \neq 1$.* (Stone [1936]) (Hint: if $x \not\sqsubseteq y$, then $\overline{x} \sqcup y \neq 1$ follows as in the proof of 20.3.7. An ultrafilter not containing $\overline{x} \sqcup y$ contains $x \sqcap \overline{y}$, and hence x but not y. Its complement is a prime ideal containing y but not x.)

[2] By 20.3.11, the Boolean Prime Ideal Theorem is equivalent to the **Boolean Maximal Ideal Theorem** saying that if $x \not\sqsubseteq y$ on a Boolean algebra, then there is a maximal ideal containing y but not x.

The Boolean Prime Ideal Theorem is equivalent to the Heyting Prime Ideal Theorem (Scott [1954a]) and hence, by note 6 on p. 422, it is not provable in ZF, it implies a weak form of the Axiom of Choice, and it does not imply the full Axiom of Choice.

b) *The Boolean Prime Ideal Theorem is equivalent to the Strong Algebraic Completeness Theorem.* (Henkin [1954]) (Hint: one direction is given by the proof of 20.2.5. Conversely, given any Boolean algebra, we can find a Lindenbaum algebra $\mathcal{A}_\Gamma$ isomorphic to it as in 20.3.1. If $x \neq 1$, then x corresponds to $[\![\alpha]\!]_\rho$, for some formula α not provable from Γ. If the Strong Algebraic Completeness Theorem holds, there is an environment ρ on $\{0, 1\}$ that sends all formulas of Γ to 1 and α to 0. Then the set $\{[\![\beta]\!]_\rho = 1\}$ is an ultrafilter of the Lindenbaum algebra not containing $[\![\alpha]\!]_\rho$, which corresponds through the isomorphism to an ultrafilter not containing x.)

Thus the use of the Boolean Prime Ideal Theorem, or of something equivalent to it, in the proof of the Strong Algebraic Completeness Theorem is not eliminable.

Clopen sets

While the Second Representation Theorem for Boolean algebras refers to power sets, and does not require any use of topology, it is unsatisfactory in that it only represents Boolean algebras as generic *sub*algebras of power sets, without characterizing them.

The consideration of Stone topologies, which were introduced to deal with $\Rightarrow$ in Heyting algebras, allows us to get an analogue of 18.4.7 for Boolean algebras, and to represent arbitrary Boolean algebras up to isomorphism.

Theorem 20.3.13 Third Representation for Boolean Algebras (Stone [1937a]) *Any Boolean algebra is isomorphic to the algebra of clopen sets of its Stone space.*

Proof. The proof of 18.4.7 shows that a Boolean algebra and the algebra of compact open sets of its Stone space are isomorphic as Heyting algebras, with an isomorphism given by the function $f : A \to \mathcal{P}(\mathcal{F}_A^p)$ defined as follows:

$$f(x) = \text{the set of all prime filters containing } x.$$

It is thus enough to show that if the Heyting algebra A is a Boolean algebra, then the set $f(A)$ of compact open sets of its Stone space coincides with the set of clopen sets:

- *every element of $f(A)$ is clopen*
  For any x, $f(x)$ is open by definition of Stone topology. Moreover, $f(\overline{x})$ is open, too. Since f preserves complements, because it preserves 0, 1, $\sqcap$ and $\sqcup$, then $\overline{f(x)} = f(\overline{x})$, i.e. $f(x)$ is closed.

- *every clopen set is in $f(A)$*
  Let X be a clopen set. Then both X and $\overline{X}$ are open, and

$$X = \bigcup_{x \in B} f(x) \qquad \text{and} \qquad \overline{X} = \bigcup_{x \in C} f(x)$$

for some subsets B and C of A, since $f(A)$ generates the Stone topology. Then

$$\mathcal{F}_A^p = X \cup \overline{X} = \bigcup_{x \in B \cup C} f(x).$$

By 18.4.8, the Stone space $\mathcal{F}_A^p$ is compact, and thus finitely many $f(x)$ with $x \in B \cup C$ are enough to cover it. In particular, there is a finite subset $\{x_1, \ldots, x_n\}$ of B such that

$$X = f(x_1) \cup \cdots \cup f(x_n).$$

Then

$$X = f(x_1 \sqcup \cdots \sqcup x_n)$$

because f preserves $\sqcup$, and thus $X \in f(A)$. □

The previous proof shows in particular that, on the Stone space of a Heyting algebra,

$$\{X : X \text{ clopen}\} \subseteq \{X : X \text{ compact open}\}.$$

Equality holds if and only if the Heyting algebra is a Boolean algebra. One direction follows from the previous proof. The other direction follows from the fact that if equality holds, then by 18.4.7 the Heyting algebra is isomorphic to an algebra of clopen sets, and hence it is a Boolean algebra.

20.4 Relationships with Heyting Algebras $\star$

The goal of the present section is to look at the relationship between the Intuitionistic and Classical Propositional Calculi from an algebraic point of view, as a relationship between Heyting and Boolean algebras. In particular, we look for a canonical way of associating a Boolean algebra to a given Heyting algebra.

Pseudocomplements

Since $\neg\alpha$ is defined as $\alpha \to \bot$ in the Intuitionistic Propositional Calculus, the following algebraic version of it in Heyting algebras will play a crucial role.

Definition 20.4.1 *In a Heyting algebra we let the* **pseudocomplement** *of an element a be the element $\sim a = a \Rightarrow 0$, i.e. the greatest element x such that $x \sqcap a = 0$.*

The next result shows that $\sim$ behaves in a Heyting algebra as $\neg$ does in the Intuitionistic Propositional Calculus, as expected.

Proposition 20.4.2 *The following laws hold in any Heyting algebra:*

1. **No Contradiction.** $a \sqcap \sim a = 0$

2. **'Good' Contrapositive.** *If* $a \sqsubseteq b$*, then* $\sim b \sqsubseteq \sim a$.

3. **'Good' Double Negation.** $a \sqsubseteq {\sim\sim}\, a$

4. **Triple Negation.** $\sim a = {\sim\sim\sim}\, a$

5. **'Good' De Morgan.** $\sim a \sqcap \sim b = \sim (a \sqcup b)$.

Proof. 1 is an obvious consequence of the definition of $\sim a$.

2 says that

$$\text{if } a \sqsubseteq b \text{ then } (b \Rightarrow 0) \sqsubseteq (a \Rightarrow 0).$$

This follows from the fact that if $a \sqsubseteq b$ and $b \sqcap x = 0$, then $a \sqcap x = 0$.

3 follows from the fact that $a \sqcap \sim a = 0$ by definition, and ${\sim\sim}\, a$ is the greatest element x such that $x \sqcap \sim a = 0$, i.e. $a \sqsubseteq {\sim\sim}\, a$.

4 follows from 3. $\sim a \sqsubseteq {\sim\sim\sim}\, a$ is a special case of it (for $\sim a$). And ${\sim\sim\sim}\, a \sqsubseteq \sim a$ follows from $a \sqsubseteq{\sim\sim}\, a$ by the 'Good' Contrapositive Law.

5 says that

$$(a \sqcup b) \Rightarrow 0 = (a \Rightarrow 0) \sqcap (b \Rightarrow 0),$$

and this follows from the fact that

$$\begin{aligned} (a \sqcup b) \sqcap x = 0 \quad &\Longleftrightarrow \quad (a \sqcap x) \sqcup (b \sqcap x) = 0 \\ &\Longleftrightarrow \quad a \sqcap x = 0 \text{ and } b \sqcap x = 0, \end{aligned}$$

by distributivity and definition of l.u.b. □

Complemented elements and clopen sets

The next result relates the two notions of complement and pseudocomplement.

Proposition 20.4.3 *If* $\overline{a}$ *exists in a Heyting algebra, then* $\overline{a} = \sim a$.

Proof. Since $\sim a$ is the greatest x such that $x \sqcap a = 0$, $\overline{a} \sqsubseteq \sim a$. From $a \sqcup \overline{a} = 1$ we get $a \sqcup \sim a = 1$. Since $a \sqcap \sim a = 0$ holds automatically, $\sim a = \overline{a}$ follows by uniqueness of complements. □

We now look at elements that are or have complements.

Definition 20.4.4 *An element* a *is* **complemented** *if* $a = \overline{b}$ *for some* b.

Notice that *a is complemented if and only if $\overline{a}$ exists*. In one direction, if $a = \overline{b}$, then $\overline{a} = b$. In the other direction, if $\overline{a}$ exists, then we can let $b = \overline{a}$.

The complemented elements provide a first way of extracting a Boolean algebra from a Heyting algebra.

Proposition 20.4.5 *Given a Heyting algebra $\mathcal{A}$, the set of complemented elements*

$$\overline{\mathcal{A}} = \{a : (\exists b)(a = \overline{b})\}$$

is a Boolean algebra.

Proof. To prove that the complemented elements form a Boolean algebra, we need the following facts:

- 0 *and* 1 *are complemented*
  It is enough to notice that $\overline{0} = 1$ and $\overline{1} = 0$, since $0 \sqcup 1 = 1$ and $0 \sqcap 1 = 0$.

- *if a and b are complemented, so are $a \sqcap b$ and $a \sqcup b$*
  By De Morgan's Laws, $\overline{a \sqcap b} = \overline{a} \sqcup \overline{b}$ and $\overline{a \sqcup b} = \overline{a} \sqcap \overline{b}$.

- *if a and b are complemented, so is $a \Rightarrow b$*
  If a is complemented, then $(a \Rightarrow b) = (\overline{a} \sqcup b)$. If b is complemented, then $\overline{a \Rightarrow b} = (a \sqcap \overline{b})$ by the De Morgan's Laws. □

Corollary 20.4.6 $\overline{\mathcal{A}}$ *is a subalgebra of* $\mathcal{A}$.

In a topological Heyting algebra,

$$\overline{A} = \text{the set-theoretical complement of } A,$$

and thus the complemented elements are the **clopen** sets.

Negative elements and regular open sets

The negative formulas, i.e. those equivalent to negations, provide a classical fragment inside the Intuitionistic Propositional Calculus (see 21.3.2). We now translate this fact into algebraic language.

Definition 20.4.7 *An element a* **negative** *if $a = \sim b$ for some b.*

Notice that *a is negative if and only if $a = \sim\sim a$.* In one direction, if $a = \sim b$, then $\sim\sim a = \sim\sim\sim b = \sim b = a$ by the Triple Negation Law. In the other direction, if $a = \sim\sim a$, then we can let $b = \sim a$.

The negative elements provide a second way of extracting a Boolean algebra from a Heyting algebra.

Proposition 20.4.8 (Tarski [1938], McKinsey and Tarski [1946], Rasiowa and Sikorski [1963]) *Given a Heyting algebra $\mathcal{A}$, the set of negative elements*

$$\sim \mathcal{A} = \{a : (\exists b)(a = \sim b)\}$$

is a Boolean algebra, and it contains $\overline{\mathcal{A}}$.

Proof. To prove that the negative elements of a Heyting algebra form a Boolean algebra, we need the following facts:

- 0 *and* 1 *are negative*
  It is enough to notice that $\sim 0 = 1$ and $\sim 1 = 0$. The former holds because $x \sqcap 0 = 0$ for every x. The latter holds because $x \sqcap 1 = 0$ only for $x = 0$.

- *if* a *and* b *are negative, then* $a \sqcap b$ *is negative (and hence it is the g.l.b. of* a *and* b *in* $\sim \mathcal{A}$*)*
  If a and b are negative, then $a = \sim a_1$ and $b = \sim b_1$. By the 'Good' De Morgan Law, $\sim a_1 \sqcap \sim b_1 = \sim (a_1 \sqcup b_1)$.

- *if* a *and* b *are negative, then* $\sim\sim (a \sqcup b)$ *is the least negative element above* $a \sqcup b$ *(and hence it is the l.u.b. of* a *and* b *in* $\sim \mathcal{A}$*)*
  It is enough to prove that, in general, $\sim\sim x$ is the least negative element above x.

  By the 'Good' Double Negation Law, $x \sqsubseteq \sim\sim x$. Conversely, suppose $x \sqsubseteq y$ and y is negative. Then $\sim\sim x \sqsubseteq \sim\sim y = y$, by a double application of the 'Good' Contrapositive Law, and because y is negative.

- *if* b *is negative, so is* $a \Rightarrow b$
  Since b is negative, $b = \sim c$ for some c. It is enough to prove

  $$(a \Rightarrow \sim c) = \sim (a \sqcap c),$$

  i.e.

  $$a \Rightarrow (c \Rightarrow 0) = (a \sqcap c) \Rightarrow 0.$$

  By definition of adjointness,

  $$\begin{aligned} x \sqsubseteq a \Rightarrow (c \Rightarrow 0) &\iff (x \sqcap a) \sqsubseteq (c \Rightarrow 0) \\ &\iff x \sqcap a \sqcap c \sqsubseteq 0 \\ &\iff x \sqsubseteq (a \sqcap c \Rightarrow 0). \end{aligned}$$

- *every negative element is complemented in* $\sim \mathcal{A}$
  We show that if a is negative, then $\sim a$ is the complement of a. Since a and $\sim a$ are both negative (the former by hypothesis, the latter by definition), by what we have already proved it is enough to show

  $$a \sqcap \sim a = 0 \qquad \text{and} \qquad \sim\sim (a \sqcup \sim a) = 1.$$

  Obviously, $a \sqcap \sim a = 0$ by definition of $\sim$. Moreover

  $$\sim (a \sqcup \sim a) = \sim a \sqcap \sim\sim a = 0$$

  by the 'Good' De Morgan Law and the definition of $\sim$. Thus $\sim\sim (a \sqcup \sim a) = 1$, because $\sim 0 = 1$.

That $\overline{\mathcal{A}} \subseteq \sim \mathcal{A}$ follows from 20.4.3, since if $a = \overline{b}$, then $a = \sim b$. □

In a topological Heyting algebra,

$$\sim A = \text{the largest open set contained in } \overline{A} = \text{the interior of } \overline{A},$$

and open sets of the form $\sim A$ are called **regular open sets**.

Exercises 20.4.9 a) $\sim \mathcal{A}$ *is the smallest subset* $\mathcal{B}$ *of* $\mathcal{A}$ *such that*:

- $0 \in \mathcal{B}$
- *if* $b \in \mathcal{B}$*, then* $a \Rightarrow b \in \mathcal{B}$.

(Hint: $\sim \mathcal{A}$ has the two properties. Moreover, if $\mathcal{B}$ has the two properties, then it contains $\sim a = (a \Rightarrow 0)$ for any $a \in \mathcal{A}$.)

b) $\sim \mathcal{A}$ *is not, in general, equal to* $\overline{\mathcal{A}}$. (Hint: consider the topological Heyting algebra of the real line. Then the open interval $(-\infty, 0)$ is negative because $(-\infty, 0) = \sim (0, \infty)$, but it is not complemented because it is not clopen.)

c) $\sim \mathcal{A}$ *is not, in general, a subalgebra of* $\mathcal{A}$. (Hint: the only way in which $\sim \mathcal{A}$ can fail to be a subalgebra of $\mathcal{A}$ is if there are negative elements a and b such that $a \sqcup b$ is not negative, since $\sim \mathcal{A}$ is closed under $\sqcap$ and $\Rightarrow$. Consider the topological algebra of part b): both the open intervals $(-\infty, 0)$ and $(0, \infty)$ are negative, but their union $\overline{\{0\}}$ is not negative, since $\sim \overline{\{0\}} = \emptyset$ while $\sim \emptyset$ is the whole space, and hence $\sim\sim \overline{\{0\}} \neq \overline{\{0\}}$.

This shows that the regular open sets of a topological space are always a Boolean algebra, but not always a field of sets.)

d) *The following laws are equivalent for a Heyting algebra*:

1. **Weak Excluded Middle.** $\sim a \sqcup \sim\sim a = 1$
2. **'Bad' De Morgan.** $\sim a \sqcup \sim b = \sim (a \sqcap b)$
3. $\sim \mathcal{A} = \overline{\mathcal{A}}$
4. $\sim \mathcal{A}$ *is a subalgebra of* $\mathcal{A}$.

(Johnstone [1979]) (Hint: for the equivalence of 1 and 2, see 21.4.6.

For the equivalence of 1 and 3, a negative element $\sim a$ is complemented in $\mathcal{A}$ if and only if $\sim a \sqcup \sim\sim a = 1$.

For the equivalence of 1 and 4, as in 21.4.5 from 1 we get filtration of $\sim\sim$ through $\sqcup$. Thus, if a and b are negative,

$$\sim\sim (a \sqcup b) = (\sim\sim a \sqcup \sim\sim b) = a \sqcup b.$$

Conversely, if $\sim A$ is a subalgebra of $\mathcal{A}$, then

$$\sim a \sqcup \sim\sim a = \sim\sim (\sim a \sqcup \sim\sim a) = \sim (\sim\sim a \sqcap \sim a) = \sim 0 = 1,$$

because $\sim a$ and $\sim\sim a$ are negative, and by the 'Good' De Morgan Law.)

Exercises 20.4.10 Completion of Boolean algebras. (MacNeille [1937], Stone [1937b]) Given a Boolean algebra $\mathcal{A}$, a complete Boolean algebra $\mathcal{C}$ is the **completion** of $\mathcal{A}$ if it contains (an isomorphic copy of) $\mathcal{A}$ as a subalgebra, and every element of $\mathcal{C}$ is the l.u.b. of a subset of $\mathcal{A}$.

a) *Every Boolean algebra has a completion.* (Hint: given a Boolean algebra $\mathcal{A}$, let $\mathcal{B}$ be the complete Heyting algebra of the open sets of its Stone space, and consider $\sim \mathcal{B}$, i.e. the complete Boolean algebra of the regular open sets. By 20.3.13, $\mathcal{A}$ is isomorphic to the algebra of clopen sets, which is a subalgebra of $\sim \mathcal{B}$. Moreover, $\sim \mathcal{B}$ is the completion of the algebra of clopen sets, since the latter generates the Stone topology.)

b) *The completion of a Boolean algebra is unique, up to isomorphism.* (Hint: given two completions $\mathcal{C}_1$ and $\mathcal{C}_2$ of $\mathcal{A}$, define an isomorphism $f : \mathcal{C}_1 \rightarrow \mathcal{C}_2$ by letting $f(x)$ be the l.u.b. in $\mathcal{C}_2$ of the set of elements of $\mathcal{A}$ below x in $\mathcal{C}_1$.)

c) *The completion of a Boolean algebra is not necessarily atomic, and in particular not necessarily isomorphic to a power set.* (Hint: if $\mathcal{A}$ is atomless, then so is its completion.)

d) *A Boolean algebra is complete if and only if it is isomorphic to the algebra of regular open sets of its Stone space.* (Hint: by part a).)

e) *A Boolean algebra is complete if and only if every regular open set of its Stone space is clopen.* (Hint: by part a) and 20.4.9.d.3.)

f) *A Boolean algebra is complete if and only if the regular open sets of its Stone space are closed under finite union.* (Hint: by part a) and 20.4.9.d.4.)

Weak units and dense open sets

By 21.1.3 the classically provable propositional formulas are exactly those whose double negation is intuitionistically provable, and the Classical Propositional Calculus can be obtained from the intuitionistic one by collapsing double negations. We now translate these facts into algebraic language.

Definition 20.4.11 *An element a is a* **weak unit** *if $\sim\sim a = 1$.*

Notice that *a is a weak unit if and only if $\sim a = 0$*, i.e. a is the only element x such that $a \sqcap x = 0$ is 0. In one direction, if $\sim\sim a = 1$, then $\sim\sim\sim a = 0$ because $\sim 1 = 0$, and $\sim a = 0$ by the Triple Negation Law. In the other direction, if $\sim a = 0$, then $\sim\sim a = 1$ because $\sim 0 = 1$.

The weak units provide a third way of extracting a Boolean algebra from a Heyting algebra.

Proposition 20.4.12 (Tarski [1938], Rasiowa and Sikorski [1963]) *Given a Heyting algebra $\mathcal{A}$, the set of weak units*

$$\{a : \sim\sim a = 1\}$$

is a filter, and the quotient Heyting algebra $\mathcal{A}_{/\sim\sim}$ generated by it is a Boolean algebra.

Proof. To prove that the weak units of a Heyting algebra form a filter, we need the following facts:

- *if a and b are weak units, so is $a \sqcap b$*
  It is enough to note that
  $$(a \sqcap b) \sqcap x = 0 \implies a \sqcap (b \sqcap x) = 0 \implies b \sqcap x = 0 \implies x = 0$$
  by distributivity, and because a and b are weak units.

- *if a is a weak unit and $a \sqsubseteq b$, then b is a weak unit*
  It is enough to note that
  $$b \sqcap x = 0 \implies a \sqcap x = 0 \implies x = 0$$
  because $a \sqsubseteq b$, and a is a weak unit.

Recall from 5.1.13 that a filter F on a Heyting algebra $\mathcal{A}$ induces an equivalence relation
$$a \sim_F b \iff (a \Rightarrow b) \in F \ \wedge \ (b \Rightarrow a) \in F,$$
and that the set of equivalence classes is a (quotient) Heyting algebra whose operations are induced by the operations of $\mathcal{A}$. It is thus enough to note that in the special case when
$$F = \{a : \sim\sim a = 1\},$$
the quotient Heyting algebra is a actually a Boolean algebra, with the complementation operation induced by $\sim$.

This is just an observation, based on the following two facts:

- $a \sqcap \sim a = 0$
  By the No Contradiction Law.

- $(a \sqcup \sim a) \in F$
  This says that $[a \sqcup \sim a]$ is the unit element of the quotient, and follows from
  $$\sim (a \sqcup \sim a) = \sim a \sqcap \sim\sim a = 0,$$
  by the 'Good' De Morgan Law and the first fact (applied to $\sim a$). □

Exercises 20.4.13 a) *$F = \{a : a = (b \sqcup \sim b)$ for some $b\}$, i.e. $\mathcal{A}_{/\sim\sim}$ is obtained by making the Excluded Middle true*. (Hint: $b \sqcup \sim b \in F$ by the proof above. Conversely, if $a \in F$, then $\sim a = 0$, i.e. $a = a \sqcup 0 = a \sqcup \sim a$.)

b) *$F = \{a : a = (\sim\sim b \Rightarrow b)$ for some $b\}$, i.e. $\mathcal{A}_{/\sim\sim}$ is obtained by making the 'Bad' Double Negation true*. (Hint: if $a \in F$, then $\sim\sim a = 1$, so $a = (1 \Rightarrow a) = (\sim\sim a \Rightarrow a)$. Conversely, $\sim\sim b \Rightarrow b \in F$ because F is upward closed, $\sim b \sqcup b \in F$ and
$$(\sim b \sqcup b) = (\sim\sim\sim b \sqcup b) \sqsubseteq (\sim\sim b \Rightarrow b),$$

where the last step follows from the fact that $(\sim c \sqcup b) \sqsubseteq (c \Rightarrow b)$.)

c) *$a \sim_F b$ if and only if $\sim\sim a = \sim\sim b$, i.e. $\mathcal{A}_{/\sim\sim}$ is obtained by identifying double negations*. (Hint: if $\sim\sim a = \sim\sim b$, then $\sim\sim a \sim_F \sim\sim b$ obviously. And $a \sim_F b$ follows from $\sim\sim a \sim_F a$ and $\sim\sim b \sim_F b$, where e.g. $\sim\sim a \sim_F a$ holds because $(a \Rightarrow \sim\sim a) = 1 \in F$, and $(\sim\sim a \Rightarrow a) \in F$ from part b).

Conversely, if $a \sim_F b$, then $a \sim_F \sim\sim b$, so $(a \Rightarrow \sim\sim b) \in F$, i.e.

$$1 = \sim\sim (a \Rightarrow \sim\sim b) = (a \Rightarrow \sim\sim b),$$

where the last equality follows from $\sim\sim (a \Rightarrow \sim c) = (a \Rightarrow \sim c)$, which we are going to prove. Then $a \sqsubseteq \sim\sim b$, and $\sim\sim a \sqsubseteq \sim\sim b$. The converse is symmetrical.

To prove $\sim\sim (a \Rightarrow \sim c) = (a \Rightarrow \sim c)$, notice that

$$(a \Rightarrow \sim c) = [a \Rightarrow (c \Rightarrow 0)] = [(a \sqcap c) \Rightarrow 0] = \sim (a \sqcap c),$$

so that

$$\sim\sim (a \Rightarrow \sim c) = \sim\sim\sim (a \sqcap c) = \sim (a \sqcap c) = (a \Rightarrow \sim c)$$

by the Triple Negation Law.)

We have seen that the negative elements of a topological Heyting algebra are the regular open sets. If we call *closure* of A the smallest closed set containing it, then

$$\sim A = \text{the complement of the closure of } A.$$

Thus the weak units of a topological Heyting algebra are the **open dense sets**, i.e. the open sets whose closure is the whole space.

The canonical Boolean algebra associated with a Heyting algebra

We have associated two Boolean algebras $\sim \mathcal{A}$ and $\mathcal{A}_{/\sim\sim}$ with a given Heyting algebra, and the next result shows that they are just different ways of looking at the same algebra.

Proposition 20.4.14 (Rasiowa and Sikorski [1963]) *Given a Heyting algebra $\mathcal{A}$, the two Boolean algebras $\sim \mathcal{A}$ and $\mathcal{A}_{/\sim\sim}$ associated with it are isomorphic.*

Proof. Let $f : \mathcal{A} \to \mathcal{A}_{/\sim\sim}$ be the canonical homomorphism of Heyting algebras defined by $f(a) = [a]$. We show that its restriction to $\sim \mathcal{A}$ is an isomorphism of Boolean algebras, as follows:

- *f preserves g.l.b.'s, the order, 0 and 1 of $\sim \mathcal{A}$*
  f preserves g.l.b.'s of $\sim \mathcal{A}$ because the latter coincide with g.l.b.'s of $\mathcal{A}$, which are preserved because f is a homomorphism. Similarly for the order, 0 and 1.

- *f preserves l.u.b.'s of $\sim\mathcal{A}$*
  If a and b are negative, their l.u.b. in $\sim\mathcal{A}$ is $\sim\sim(a \sqcup b)$. Since f preserves l.u.b.'s of $\mathcal{A}$, it is enough to show that $[a \sqcup b] = [\sim\sim(a \sqcup b)]$.

  We show more generally that $[x] = [\sim\sim x]$, which follows by proving that both $x \Rightarrow \sim\sim x$ and $\sim\sim x \Rightarrow x$ are in the filter F of weak units.

  Since $x \sqsubseteq \sim\sim x$ we have

  $$1 = (x \Rightarrow x) \sqsubseteq (x \Rightarrow \sim\sim x)$$

  by 5.1.8.7 and monotonicity of $\Rightarrow$, and hence $x \Rightarrow \sim\sim x$ is in any filter.

  Since we know that $(x \sqcup \sim x) \in F$, to get that $(\sim\sim x \Rightarrow x) \in F$ it is enough to show that $(x \sqcup \sim x) \sqsubseteq (\sim\sim x \Rightarrow x)$. By the Triple Negation Law, it is equivalent to show that $(x \sqcup \sim\sim\sim x) \sqsubseteq (\sim\sim x \Rightarrow x)$, which follows from $(x \sqcup \sim y) \sqsubseteq (y \Rightarrow x)$, by taking $y = \sim\sim x$.

  Finally, $(x \sqcup \sim y) \sqsubseteq (y \Rightarrow x)$ holds as follows:

  $$\begin{array}{rl} & (x \sqcup \sim y) \sqsubseteq (y \Rightarrow x) \\ \Longleftrightarrow & (x \sqcup \sim y) \sqcap y \sqsubseteq x \\ \Longleftrightarrow & (x \sqcap y) \sqcup (y \sqcap \sim y) \sqsubseteq x \\ \Longleftrightarrow & (x \sqcap y) \sqsubseteq x, \end{array}$$

  by adjointness, distributivity and the fact that $y \sqcap \sim y = 0$.

- *f is one-one on $\sim\mathcal{A}$*
  We want to show that if a and b are negative and $[a] \sqsubseteq [b]$, i.e. $(a \Rightarrow b) \in F$, then $a \sqsubseteq b$.

  The hypotheses are equivalent to $\sim\sim a = a$, $\sim\sim b = b$ and $\sim\sim(a \Rightarrow b) = 1$. Then $(\sim\sim a \Rightarrow \sim\sim b) = 1$ by the last hypothesis and the 'Good' Contrapositive Law (applied twice), and $(a \Rightarrow b) = 1$ by the first two hypotheses. This implies $a \sqsubseteq b$, because by adjointness

  $$1 \sqsubseteq a \Rightarrow b \iff 1 \sqcap a \sqsubseteq b \iff a \sqsubseteq b.$$

- *the restriction of f to $\sim\mathcal{A}$ is onto $\mathcal{A}_{/\sim\sim}$*
  We proved above that $[a] = [\sim\sim a]$. Then every element of $\sim\mathcal{A}$ is the image via f of a negative element. □

Conditions for a Heyting algebra to be a Boolean algebra

The previous results provide a satisfactory answer to the question of how to associate a canonical Boolean algebra to a given Heyting algebra, and we now turn to the complementary question of when is a given Heyting algebra already a Boolean algebra.

The first answer we provide is *local algebraic*, and it translates the fact that the Classical Propositional Calculus is obtained from the Intuitionistic Propositional Calculus by adding either the Law of the Excluded Middle or the Double Negation Law (see 21.2.3). It also shows that the pseudocomplements of a Boolean algebra are simply the complements.

Proposition 20.4.15 *The following laws are equivalent for a Heyting algebra:*

1. **Excluded Middle.** $a \sqcup \sim a = 1$
2. **'Bad' Double Negation.** $\sim\sim a \sqsubseteq a$
3. **'Bad' Contrapositive.** *If* $\sim a \sqsubseteq \sim b$ *then* $b \sqsubseteq a$*,*

and a Heyting algebra is a Boolean algebra if and only if it satisfies (any of) them.

Proof. To prove the equivalence of 1 and 2, suppose $a \sqcup \sim a = 1$. Then

$$\sim\sim a = \sim\sim a \sqcap 1 = \sim\sim a \sqcap (a \sqcup \sim a) = \ (\sim\sim a \sqcap a) \sqcup (\sim\sim a \sqcap \sim a) = \sim\sim a \sqcap a$$

by distributivity, and the fact that $\sim\sim a \sqcap \sim a = 0$. Then $\sim\sim a \sqsubseteq a$.

Conversely, suppose $\sim\sim (a \sqcup \sim a) \sqsubseteq (a \sqcup \sim a)$. Then $a \sqcup \sim a = 1$, because

$$\sim\sim (a \sqcup \sim a) = \sim (\sim a \sqcap \sim\sim a) = \sim 0 = 1$$

by the 'Good' De Morgan Law.

To prove the equivalence of 2 and 3, suppose $\sim a \sqsubseteq \sim b$. Then $\sim\sim b \sqsubseteq \sim\sim a$ by the 'Good' Contrapositive Law, and

$$b \sqsubseteq \sim\sim b \sqsubseteq \sim\sim a \sqsubseteq a$$

by the 'Good' and 'Bad' Double Negation Laws.

Conversely, $\sim a \sqsubseteq \sim\sim\sim a$ by the 'Good' Double Negation Law applied to $\sim a$, so $\sim\sim a \sqsubseteq a$ by 3.

Because of the equivalences just proved, it is now enough to show that a Heyting algebra is a Boolean algebra if and only if it satisfies the Law of the Excluded Middle. In one direction, since $a \sqcap \sim a = 0$ always holds, the Law implies that every element is complemented (with $\sim$ as the complementation operation). In the other direction, in a Boolean algebra $\sim$ is the complementation operation, and thus the Law is satified. □

Exercise 20.4.16 *A Heyting algebra is a Boolean algebra if and only if* 1 *is the only weak unit.* (Hint: if $\sim\sim a \sqsubseteq a$, then from $\sim\sim a = 1$ we have $a = 1$. Conversely, since

$\sim\sim (a \sqcup \sim a) = 1$ always holds, if 1 is the only weak unit, then $a \sqcup \sim a = 1$.)

The second answer to the problem stated at the beginning of the subsection is *global algebraic*.

Proposition 20.4.17 (Nachbin [1947]) *A Heyting algebra is a Boolean algebra if and only if every prime filter is maximal.*

Proof. We already know from 20.3.11 that in a Boolean algebra every prime filter is maximal.

For the converse, suppose that a Heyting algebra has a non complemented element a. Consider first the filter J of all elements joining to 1 with a, i.e.

$$J = \{x : a \sqcup x = 1\},$$

and then the filter F generated by $J \cup \{a\}$. By distributivity, F is the upward closure of

$$\{a \sqcap x : x \in J\} = \{a \sqcap x : a \sqcup x = 1\}.$$

Then F is a proper filter, otherwise $0 \in F$, and there would be an x such that

$$0 = a \sqcap x \qquad \text{and} \qquad a \sqcup x = 1,$$

i.e. a would be complemented.

We want to prove that there is a prime filter S contained in F and avoiding a, so that $S \subset F$ and S is not maximal. Consider the ideal I generated by $\overline{F} \cup \{a\}$. It is disjoint from J, otherwise there is $x \in J$ (i.e. satisfying $a \sqcup x = 1$) such that $x \sqsubseteq a \sqcup y$ for some $y \in \overline{F}$. But since $a \sqcup x = 1$, $a \sqcup y = 1$ too, and $y \in J \subseteq F$, contradiction.

We now prove that there is a prime filter S extending J and disjoint from I: in particular, S is a prime filter contained in F and avoiding a. We consider the family $\mathcal{F}_A$ of all filters extending J and disjoint from I, and notice that since every chain in this set has a l.u.b. (namely, its union), by Zorn's Lemma there is a maximal element S. It only remains to prove that S is prime.

Suppose S is not prime. Then there are $x, y \notin S$ such that $x \sqcup y \in S$, and the filters generated by $S \cup \{x\}$ and $S \cup \{y\}$ are not in the family $\mathcal{F}_A$, because they properly contain the maximal element S. Since the two filters contain J by definition, this means that they are not disjoint from I. Then there must exist elements

$$\begin{array}{lll} a_1 \in I \text{ and } b_1 \in S & \text{such that} & a_1 \sqsupseteq x \sqcap b_1 \\ a_2 \in I \text{ and } b_2 \in S & \text{such that} & a_2 \sqsupseteq y \sqcap b_2. \end{array}$$

Then

$$a_1 \sqcup a_2 \sqsupseteq (x \sqcap b_1) \sqcup (y \sqcap b_2) = (x \sqcup y) \sqcap (b_1 \sqcup y) \sqcap (x \sqcup b_2) \sqcap (b_1 \sqcup b_2),$$

by distributivity. The last expression is in S because the first conjunct is in S by hypothesis, and the remaining three are in S because each is above one of b_1 and b_2, and S is a filter. By upward closure of S, then $a_1 \sqcup a_2 \in S$ too, and hence $S \cap I \neq \emptyset$, contradicting the fact that S is disjoint from I. □

The third and final answer to the problem stated at the beginning of the subsection is *topological*. Recall from 18.3.5 that a topology is T_2 if any pair of distinct elements can be separated by disjoint open sets.

Proposition 20.4.18 (Stone [1937]) *A Heyting algebra is a Boolean algebra if and only if its Stone topology is T_2.*

Proof. Let f be the function from a Boolean algebra A to its Stone space $\mathcal{F}^p_A$ considered in 20.3.2 and 20.3.13, i.e.

$$f(x) = \text{the set of all prime filters containing } x.$$

Since f is a homomorphism of Boolean algebras, if $x \sqcap y = 0$, then $f(x) \cap f(y) = \emptyset$, i.e. $f(x)$ and $f(y)$ are disjoint open sets.

To show that $\mathcal{F}^p_A$ is T_2 it is thus enough to find, for any pair F and G of distinct elements of $\mathcal{F}^p_A$ (i.e. distinct prime filters on A) two elements x and y of A such that $x \sqcap y = 0$, $F \in f(x)$ and $G \in f(y)$.

Since F and G are distinct, there is an element x on which they differ, e.g. $x \in F - G$. Then $\sim x \in G - F$ (if $x \in F$, then $\sim x \notin F$, otherwise $x \sqcap \sim x = 0 \in F$; if $x \notin G$, then $\sim x \in G$, because a prime filter is an ultrafilter). Thus $x \sqcap \sim x = 0$, $F \in f(x)$ because $x \in F$, and $G \in f(\sim x)$ because $\sim x \in G$.

Conversely, let A be a Heyting algebra. By 18.4.7, A is isomorphic to the algebra of compact open sets of its Stone space, and this algebra contains all clopen sets. We show that if the Stone space is T_2, then every compact open set is closed, so that A is isomorphic to an algebra of clopen sets, and in particular it is a Boolean algebra.

Let X be a compact open set in the Stone space, and $y \notin X$. We want to find an open set O_y containing y and disjoint from X, so that $\overline{X} = \bigcup_{y \in X} O_y$. Thus $\overline{X}$ is open, and X is closed.

If $x \in X$, then $x \neq y$ because $y \notin X$. Since the Stone space is T_2, there are disjoint open sets B_x and C_x such that $x \in B_x$ and $y \in C_x$. Since the Stone space is compact (by 18.4.8) and $X \subseteq \bigcup_{x \in X} A_x$, there are $x_1, \ldots, x_n$ such that $X \subseteq A_{x_1} \cup \cdots \cup A_{x_n}$. Then $O_y = B_{x_1} \cap \cdots \cap B_{x_n}$ is an open set containing y and disjoint from X, as needed. □

Exercise 20.4.19 *A topology is homeomorphic to the dual Stone topology of a Boolean algebra if and only if it is coherent and T_2.* (Stone [1937]) (Hint: by 18.5.7 and 20.4.18.)

20.5 Boolean Bicartesian Closed Categories ⋆

Having introduced a classical analogue of Heyting algebras, we now look for a classical analogue of bicartesian closed categories. According to the motto

$$\text{bicartesian closed categories} \ = \ \text{Heyting algebras} + \text{morphisms},$$

the following is the obvious notion.

Definition 20.5.1 *A bicartesian closed category is called* **Boolean** *if the underlining Heyting algebra is a Boolean algebra.*

The previous motto now becomes:

$$\text{Boolean bicartesian closed categories} \ = \ \text{Boolean algebras} + \text{morphisms}.$$

The next result shows that, unlike the intuitionistic case, the classical case collapses.

Proposition 20.5.2 (Joyal) *A Boolean bicartesian closed category is a Boolean algebra.*

Proof. We need to show that, given two objects A and B, there is at most one morphism between them. Since the category is Boolean, and hence the underlining algebra satisfies the double negation law, the object B is isomorphic to $\sim\sim B$, and hence to an object $C \Rightarrow 0$ (where C is $B \Rightarrow 0$). By cartesian closure,

$$\mathit{Hom}(A, B) \cong \mathit{Hom}(A, C \Rightarrow 0) \cong \mathit{Hom}(A \times C, 0).$$

If we show that in a cartesian closed category with 0 the set $\mathit{Hom}(X, 0)$ has at most one element, then the right-hand-side has at most one element, and hence so does the left-hand-side.

We prove that if $\mathit{Hom}(X, 0)$ is not empty then $X \cong 0$: thus either $\mathit{Hom}(X, 0)$ is empty or it can have only one element, by definition of 0.

Suppose $\mathit{Hom}(A, 0) \neq \emptyset$: then there is at least one morphism $h : A \to 0$. Since 0 is initial, there is exactly one morphism $i_X : 0 \to X$. To prove that $X \cong 0$, it is enough to show the following:

- $h \circ i_X = id_0$

  Both $h \circ i_X$ and id_0 are morphisms from 0 to the same object (again 0): by definition of initial object there can be only one such morphism, so they must be equal.

- $i_X \circ h = id_X$

  By definition of product $0 \times X$

$$h = l_{0,X} \circ \langle h, id_X \rangle \qquad \text{and} \qquad i_X = r_{0,X} \circ \langle id_0, i_X \rangle,$$

and hence

$$i_X \circ h = r_{0,X} \circ \langle id_0, i_X \rangle \circ l_{0,X} \circ \langle h, id_X \rangle,$$

i.e.

$$X \xrightarrow{\langle h, id_X \rangle} 0 \times X \xrightarrow{l_{0,X}} 0 \xrightarrow{\langle id_0, i_X \rangle} 0 \times X \xrightarrow{r_{0,X}} X.$$

Let us consider the central morphism

$$\langle id_0, i_X \rangle \circ l_{0,X} \in Hom(0 \times X, 0 \times X):$$

by cartesian closure

$$Hom(0 \times X, 0 \times X) \cong Hom(0, (0 \times X)^X),$$

and since the right-hand-side contains only one element because 0 is initial, so does the left-hand-side; since obviously

$$id_{0 \times X} \in Hom(0 \times X, 0 \times X),$$

it follows that

$$\langle id_0, i_X \rangle \circ l_{0,X} = id_{0 \times X}.$$

Then

$$\begin{aligned} i_X \circ h &= r_{0,X} \circ id_{0 \times X} \circ \langle h, id_X \rangle \\ &= r_{0,X} \circ \langle h, id_X \rangle \\ &= id_X, \end{aligned}$$

by the properties of the identity $id_{0 \times X}$, and of the projection $r_{0,X}$. □

The result just proved can be rephrased in the following proportion:

$$\frac{\text{Heyting algebras}}{\text{bicartesian closed categories}} = \frac{\text{Boolean algebras}}{\text{Boolean algebras,}}$$

and shows that there is no categorical analogue of Boolean algebras. In particular, category theory is inherently intuitionistic, and it could not have been anticipated in a simplified form in a classical framework (in the same way as Heyting algebras were anticipated by Boolean algebras).

More generally, the proof of the result also sheds light on the proof theory of negative formulas (possibly from premises), by showing that there can be at most one such proof for each such formula, up to equivalence. In other words, if such formulas are provable at all (it does not matter whether classically or intuitionistically, by 21.3.2), then they are provable in an essentially unique way.

æ

Chapter 21

Relationships with Intuitionism

In this chapter we consider various relationships between the Classical and the Intuitionistic Propositional Calculi, from two points of view. First of all, we find extensions of the intuitionistic systems $\mathcal{N}$ and $\mathcal{H}$ by new axioms, thus providing approaches to the Classical Propositional Calculus in the styles of Natural Deduction and Hilbert systems. Secondly, we look for syntactical transformations of formulas that allow translations of the Classical Propositional Calculus into the intuitionistic one. The two approaches are complementary, and they respectively expand the intuitionistic notion of proof, and the classical interpretation of formulas.

21.1 Intuitionistic Analysis of Tautologies

We introduce the main technical tool of the chapter, namely an analysis of the notion of tautology from an intuitionistic point of view.

Excluded Middle

We start by showing that the laws governing the formation of classical truth-tables can actually be proved in the intuitionistic system $\mathcal{N}$, if we identify truth and falsity assertions about a formula with, respectively, the formula itself and its negation.

We consider all propositional connectives, although classically $\neg$ and $\wedge$ would suffice (see 19.2), because we want to analize different translations of the Classical Propositional Calculus into the intuitionistic one, while the restriction to $\neg$ and $\wedge$ would instead correspond to a fixed translation (see 21.3.1).

Proposition 21.1.1 *The formation laws of classical truth-tables are provable in $\mathcal{N}$. More precisely, for every α and β:*

1. **Negation.** *$\neg\alpha$ is true if α is false, and false if α is true:*

$$\neg\alpha \vdash_{\mathcal{N}} \neg\alpha \qquad \alpha \vdash_{\mathcal{N}} \neg\neg\alpha.$$

2. **Conjunction.** *$\alpha \wedge \beta$ is true if both α and β are true, and false if one of them is false:*

$$\alpha, \beta \vdash_{\mathcal{N}} \alpha \wedge \beta \qquad \neg\alpha \vdash_{\mathcal{N}} \neg(\alpha \wedge \beta) \qquad \neg\beta \vdash_{\mathcal{N}} \neg(\alpha \wedge \beta).$$

3. **Disjunction.** *$\alpha \vee \beta$ is true if one of α and β is true, and false if they are both false:*

$$\alpha \vdash_{\mathcal{N}} \alpha \vee \beta \qquad \beta \vdash_{\mathcal{N}} \alpha \vee \beta \qquad \neg\alpha, \neg\beta \vdash_{\mathcal{N}} \neg(\alpha \vee \beta).$$

4. **Implication.** *$\alpha \to \beta$ is true if α is false or β is true, and false if α is true and β is false:*

$$\neg\alpha \vdash_{\mathcal{N}} \alpha \to \beta \qquad \beta \vdash_{\mathcal{N}} \alpha \to \beta \qquad \alpha, \neg\beta \vdash_{\mathcal{N}} \neg(\alpha \to \beta).$$

Proof. The first property of negation is an axiom, and the second is proved by the following:

$$\dfrac{\dfrac{\dfrac{[\alpha]^{(2)} \qquad [\neg\alpha]^{(1)}}{\bot}}{(\neg\neg\alpha)^{(1)}}}{\alpha^{(2)} \to \neg\neg\alpha,}$$

The first property of conjunction is the rule of ∧-introduction. The second is proved by the following:

$$\dfrac{\dfrac{\dfrac{[\alpha \wedge \beta]^{(1)}}{\alpha} \qquad \neg\alpha}{\bot}}{\neg(\alpha \wedge \beta)^{(1)},}$$

and similarly for the third property.

The first two properties of disjunction are the rules of ∨-introduction, and the last one is proved by the following:

$$\dfrac{\dfrac{[\alpha \vee \beta]^{(2)} \qquad \dfrac{[\alpha]^{(1)} \qquad \neg\alpha}{\bot} \qquad \dfrac{[\beta]^{(1)} \qquad \neg\beta}{\bot}}{\bot^{(1)}}}{\neg(\alpha \vee \beta)^{(2)}.}$$

The first property of implication is proved by the following (using $\bot$-elimination):

$$\dfrac{\dfrac{\dfrac{[\alpha]^{(1)} \qquad \neg\alpha}{\bot}}{\beta}}{\alpha^{(1)} \to \beta.}$$

The second property follows from the rule of $\to$-introduction, and the last one is proved by the following:

$$\dfrac{\dfrac{\dfrac{\alpha \qquad [\alpha \to \beta]^{(1)}}{\beta} \qquad \neg\beta}{\bot}}{\neg(\alpha \to \beta)^{(1)}.} \qquad \square$$

Knowing that the formation rules of truth-tables can be proved in $\mathcal{N}$ we immediately have, by induction on the construction of a truth-table, that also any row of any truth-table can be so proved.

Theorem 21.1.2 Intuitionistic Provability of Tautologies from the Excluded Middle (Gentzen [1934]) *For any tautology α, if the letters occurring in it are among $p_1, \ldots, p_n$, then*

$$p_1 \vee \neg p_1, \ldots, p_n \vee \neg p_n \vdash_{\mathcal{N}} \alpha.$$

Proof. By iterated $\vee$-eliminations, to prove

$$p_1 \vee \neg p_1, \ldots, p_n \vee \neg p_n \vdash_{\mathcal{N}} \alpha$$

it is enough to prove α from all possible 2^n sequences $\overline{p}_1, \ldots, \overline{p}_n$, in which each $\overline{p}_i$ is one of p_i and $\neg p_i$.

If, as above, we identify truth and falsity assertions about a formula with, respectively, the formula itself and its negation, such sequences $\overline{p}_1, \ldots, \overline{p}_n$ can be identified with a truth assignments to the letters $p_1, \ldots, p_n$, where p_i is true or false in the given assignment according to whether $\overline{p}_i$ is either p_i or $\neg p_i$.

Then the fact that α is a tautology, i.e. that any row of a truth-table for α makes it true, implies by induction on the construction of the truth-table (using 21.1.1) that α can be proved from any such sequence $\overline{p}_1, \ldots, \overline{p}_n$. $\square$

Double Negation

In the previous subsection we kept tautologies as they are, and accounted for them in the Intuitionistic Propositional Calculus by adding axioms to it. In the present subsection we keep the Intuitionistic Propositional Calculus as it is, and account for the tautologies by doubly negating them.

The next result provides an intuitionistic justification of the non intuitionistic principles of the Classical Propositional Calculus, such as the Law of the Excluded Middle. It shows that, although these principles cannot be proved outright, at least their negation can be refuted. This strategy was first used by Aristotle (*Metaphysics*, Γ, 1006a) in justifying the unprovable principles of methaphysics, including the Law of the Excluded Middle.

Theorem 21.1.3 Double Negation of Tautologies (Kolmogorov [1925], Glivenko [1929], Gödel [1933], Gentzen [1933]) *For any α,*

$$\models \alpha \;\Leftrightarrow \vdash_{\mathcal{N}} \neg\neg\alpha.$$

Proof. We only have to prove the left to right direction, since α and $\neg\neg\alpha$ are classically equivalent.

By the $\wedge$ and $\rightarrow$ introduction rules, in 21.1.2 we have actually proved that if α is a tautology, then:

$$\vdash_{\mathcal{N}} [\bigwedge_{1\leq i\leq n} (p_i \vee \neg p_i)] \rightarrow \alpha.$$

We prove below the following *'Good' Contrapositive Law*:

$$\vdash_{\mathcal{N}} (\beta \rightarrow \gamma) \rightarrow (\neg\gamma \rightarrow \neg\beta).$$

By using it twice we then get

$$\vdash_{\mathcal{N}} \neg\neg[\bigwedge_{1\leq i\leq n} (p_i \vee \neg p_i)] \rightarrow \neg\neg\alpha.$$

We then show that double negation filters through finite conjunctions, so that

$$\vdash_{\mathcal{N}} [\bigwedge_{1\leq i\leq n} \neg\neg(p_i \vee \neg p_i)] \rightarrow \neg\neg\alpha.$$

Finally we notice that, unlike $p \vee \neg p$ itself, $\neg\neg(p \vee \neg p)$ is actually provable in $\mathcal{N}$ for any p, and thus $\neg\neg\alpha$ is too.

1. *'Good' Contrapositive Law*
   We prove that, for any β and γ,

   $$\vdash_{\mathcal{N}} (\beta \rightarrow \gamma) \rightarrow (\neg\gamma \rightarrow \neg\beta).$$

   Since negation is defined in terms of implication, it is enough to prove the *Transitive Law of Implication*

   $$(\beta \rightarrow \gamma) \rightarrow [(\gamma \rightarrow \delta) \rightarrow (\beta \rightarrow \delta)],$$

from which the claim can be obtained by letting $\delta = \bot$. Transitivity is proved by the following:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{[\beta \to \gamma]^{(3)} \qquad [\beta]^{(1)}}{\gamma} \qquad [\gamma \to \delta]^{(2)}}{\delta}}{\beta^{(1)} \to \delta}}{(\gamma \to \delta)^{(2)} \to (\beta \to \delta)}}{(\beta \to \gamma)^{(3)} \to [(\gamma \to \delta) \to (\beta \to \delta)].}$$

2. *double negation filters through conjunction*
The argument given at the beginning requires a step from

$$\vdash_{\mathcal{N}} \neg\neg(\beta \wedge \gamma) \;\to\; \neg\neg\alpha,$$

where β and γ are instances of the excluded middle, to

$$\vdash_{\mathcal{N}} (\neg\neg\beta \;\wedge\; \neg\neg\gamma) \;\to\; \neg\neg\alpha.$$

For this, we only need the following:

$$\vdash_{\mathcal{N}} (\neg\neg\beta \wedge \neg\neg\gamma) \to \neg\neg(\beta \wedge \gamma)$$

or, by the $\wedge$ and $\to$ introduction rules,

$$\neg\neg\beta, \neg\neg\gamma \vdash_{\mathcal{N}} \neg\neg(\beta \wedge \gamma).$$

This is proved by the following:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{[\neg(\beta \wedge \gamma)]^{(3)} \qquad \dfrac{[\beta]^{(1)} \qquad [\gamma]^{(2)}}{\beta \wedge \gamma}}{\bot}}{\neg\beta^{(1)}} \qquad \neg\neg\beta}{\bot}}{\neg\gamma^{(2)}} \qquad \neg\neg\gamma}{\bot}}{\neg\neg(\beta \wedge \gamma)^{(3)}.}$$

Although we don't need it in the present proof, we notice for future use that also the opposite implication holds. Indeed, from $\beta \wedge \gamma \to \beta$ we have, by a double application of the 'Good' Contrapositive Law, $\neg\neg(\beta \wedge \gamma) \to \neg\neg\beta$, and similarly for γ, i.e.

$$\vdash_{\mathcal{N}} \neg\neg(\beta \wedge \gamma) \to \neg\neg\beta \wedge \neg\neg\gamma.$$

3. *double negation of the excluded middle*
   This is proved by the following:

$$\dfrac{\dfrac{[\neg(p \vee \neg p)]^{(2)} \qquad \dfrac{\dfrac{\dfrac{[\neg(p \vee \neg p)]^{(2)} \qquad \dfrac{[p]^{(1)}}{p \vee \neg p}}{\bot}}{\neg p^{(1)}}}{p \vee \neg p}}{\bot}}{\neg\neg(p \vee \neg p)^{(2)}.}$$

This finishes the proof. □

Notice that the proof can be seen as a reduction of the double negation of tautologies to double negations of atomic instances of the Law of the Excluded Middle, which are then proved outright.

Exercise 21.1.4 *Prove the double negation of the Law of the Excluded Middle by forcing.* (Hint: trivial, by definition of forcing for negation.)

This proof is much more direct than the one in $\mathcal{N}$, but it does not give any hint on how to find the latter. A better hint is provided by a proof in $\mathcal{S}$.

While the previous proof only provides a translation of classical *validity*, a small addition provides a similar translation of classical *logical consequence*.

Corollary 21.1.5 The Double Negation Translation. *For every* Γ *and* α*,*

$$\Gamma \models \alpha \;\Leftrightarrow\; \neg\neg\Gamma \vdash_{\mathcal{N}} \neg\neg\alpha,$$

where $\neg\neg\Gamma = \{\neg\neg\gamma : \gamma \in \Gamma\}$.

Proof. We first prove the following:

- *double negation filters through implication*
  We prove that

$$\vdash_{\mathcal{N}} \neg\neg(\alpha \rightarrow \beta) \leftrightarrow (\neg\neg\alpha \rightarrow \neg\neg\beta).$$

The left to right direction follows from:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{[\alpha]^{(1)} \quad [\alpha \to \beta]^{(2)}}{\beta} \quad [\neg\beta]^{(3)}}{\bot}}{\neg\alpha^{(1)}} \quad [\neg\neg\alpha]^{(4)}}{\bot}}{\neg(\alpha \to \beta)^{(2)}} \quad [\neg\neg(\alpha \to \beta)]^{(5)}}{\bot}}{\neg\neg\beta^{(3)}}}{\neg\neg\alpha^{(4)} \to \neg\neg\beta}}{\neg\neg(\alpha \to \beta)^{(5)} \to (\neg\neg\alpha \to \neg\neg\beta).}$$

The right to left direction follows from the following, by using $\bot$-elimination:

$$\dfrac{\dfrac{\dfrac{\dfrac{[\neg\neg\alpha \to \neg\neg\beta]^{(5)} \quad \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{[\neg\alpha]^{(2)} \quad [\alpha]^{(1)}}{\bot}}{\beta}}{\alpha^{(1)} \to \beta} \quad [\neg(\alpha \to \beta)]^{(4)}}{\bot}}{\neg\neg\alpha^{(2)}}}{\neg\neg\beta} \quad \dfrac{\dfrac{\dfrac{[\beta]^{(3)}}{\alpha \to \beta} \quad [\neg(\alpha \to \beta)]^{(4)}}{\bot}}{\neg\beta^{(3)}}}{\bot}}{\neg\neg(\alpha \to \beta)^{(4)}}}{(\neg\neg\alpha \to \neg\neg\beta)^{(5)} \to \neg\neg(\alpha \to \beta).}$$

To prove the corollary, by the Compactness Theorem we restrict our attention to finite sets of formulas $\{\gamma_1, \dots, \gamma_n\}$. Then:

$$\begin{array}{rl} & \gamma_1, \dots, \gamma_n \models \alpha \\ \Leftrightarrow & \models \gamma_1 \wedge \cdots \wedge \gamma_n \to \alpha \\ \Leftrightarrow & \vdash_{\mathcal{N}} \neg\neg(\gamma_1 \wedge \cdots \wedge \gamma_n \to \alpha) \\ \Leftrightarrow & \vdash_{\mathcal{N}} \neg\neg\gamma_1 \wedge \cdots \wedge \neg\neg\gamma_n \to \neg\neg\alpha \\ \Leftrightarrow & \neg\neg\gamma_1, \dots, \neg\neg\gamma_n \vdash_{\mathcal{N}} \neg\neg\alpha \end{array}$$

by the Deduction Theorem, 21.1.3, the fact that double negation filters through both $\wedge$ and $\to$ (as proved in 21.1.3 and above), and $\to$ and $\wedge$ introduction. □

Exercise 21.1.6 *Prove the Double Negation Translation in the form*

$$\Gamma \models \alpha \;\Leftrightarrow\; \neg\neg\Gamma \vdash_{\mathcal{H}} \neg\neg\alpha.$$

(Hint: if $\Gamma \models \alpha$, then $\Gamma \vdash_{\mathcal{HC}} \alpha$ by 21.1.2, where $\mathcal{HC}$ is obtained from $\mathcal{H}$ by adding the Law of the Excluded Middle as a schema of axioms. We can then take a proof of α from Γ in $\mathcal{HC}$, and insert $\neg\neg$ everywhere. To turn the result into a proof of $\neg\neg\alpha$ from $\neg\neg\Gamma$ in $\mathcal{H}$, we then only need to provide the following in $\mathcal{H}$:

- for each instance $p \vee \neg p$ of the Law of the Excluded Middle, a proof of $\neg\neg(p \vee \neg\neg p)$
- for each axiom β of $\mathcal{H}$, a proof of $\neg\neg\beta$ from β (which can be used as an axiom) and $\beta \rightarrow \neg\neg\beta$ (which is provable in $\mathcal{H}$)
- for each application of Modus Ponens

$$\frac{\gamma \quad \gamma \rightarrow \delta}{\delta,}$$

  a proof of $\neg\neg\delta$ from $\neg\neg\gamma$ and $\neg\neg\gamma \rightarrow \neg\neg\delta$. The latter can be deduced from $\neg\neg(\gamma \rightarrow \delta)$ by a double application of the 'Good' Contrapositive Law).

Translations such as the ones discussed above are interesting because they provide *relative consistency proofs*: if it is possible to prove a contradiction (namely, both α and $\neg\alpha$) in the Classical Propositional Calculus, then it is possible to prove a contradiction in the Intuitionistic Propositional Calculus (namely $\neg\neg\alpha$ and $\neg\alpha$). This is of little interest in the context of propositional logic, where consistency is provable by finitary means, but it becomes interesting with stronger systems, where consistency becomes problematic.

21.2 Extensions of the Intuitionistic Propositional Calculus

The results of the previous section suggest the possibility of seeing the Classical Propositional Calculus as an extension of the intuitionistic one, obtained by adding to the latter the Laws of the Excluded Middle or of Double Negation. We now spell out this suggestion, thus obtaining a number of formulations of the Classical Propositional Calculus.

Excluded Middle and Double Negation again

First we notice that the two laws have the same strength from the intuitionistic point of view.

Proposition 21.2.1 *The following two laws:*

1. **Excluded Middle.** $\alpha \vee \neg\alpha$
2. **'Bad' Double Negation.** $\neg\neg\alpha \rightarrow \alpha$

are intuitionistically equivalent and not valid.

Proof. 1 implies 2 as follows:

$$\dfrac{\dfrac{\alpha \vee \neg\alpha \qquad [\alpha]^{(1)} \qquad \dfrac{\dfrac{[\neg\alpha]^{(1)} \qquad [\neg\neg\alpha]^{(2)}}{\bot}}{\alpha}}{\alpha^{(1)}}}{\neg\neg\alpha^{(2)} \to \alpha.}$$

Conversely, 2 implies 1 as follows. By assuming 2 for any formula, we have

$$\neg\neg(\alpha \vee \neg\alpha) \to (\alpha \vee \neg\alpha).$$

By 21.1.3 the premise is intuitionistically provable, because $\alpha \vee \neg\alpha$ is a classical tautology. The conclusion then follows.

Because of the equivalence just proved, it is enough to show that the Law of the Excluded Middle is not intuitionistically valid. This follows from the fact that it fails in a Kripke model with two nodes $\emptyset$ and 0, and such that $\mathcal{A}_\emptyset = \{\emptyset\}$, and $\mathcal{A}_0 = \{\alpha\}$. □

The intuitionistic equivalence of the two *schemata* cannot be improved to an equivalence of *formulas*. Indeed, while

$$(\alpha \vee \neg\alpha) \to (\neg\neg\alpha \to \alpha)$$

has been proved in the proof above, the converse implication

$$(\neg\neg\alpha \to \alpha) \to (\alpha \vee \neg\alpha)$$

fails intuitionistically, because it fails in a Kripke model with three nodes $\emptyset$, 0 and 1, and such that $\mathcal{A}_\emptyset = \mathcal{A}_0 = \emptyset$ and $\mathcal{A}_1 = \{\alpha\}$.

A number of intuitionistic principles turn out to be equivalent to the two just discussed, and the next exercise provides an additional one.

Exercise 21.2.2 The 'Bad' Contrapositive Law. We have seen in the proof of 21.1.3 that the 'Good' Contrapositive Law is intuitionistically valid. We now consider the 'bad' one, namely

$$(\neg\alpha \to \neg\beta) \to (\beta \to \alpha).$$

The 'Bad' Contrapositive Law is equivalent to the Law of the Excluded Middle. (Hint: it is easier to prove the equivalence with the 'Bad' Double Negation Law. In one direction, we use the 'Good' Contrapositive Law and then eliminate double negations by transitivity of implication. In the other direction, we let $\beta = \neg\neg\alpha$ in the 'Bad' Contrapositive Law and notice that the premise is intuitionistically valid.)

Both the Law of the Excluded Middle and the 'Bad' Double Negation Law are stated for arbitrary formulas, but for the purpose of axiomatizing the Classical Propositional Calculus weaker schemata (namely, only the atomic instances of any of the two laws) are enough.

Proposition 21.2.3 *The Classical Propositional Calculus can be obtained by adding to either $\mathcal{H}$ or $\mathcal{N}$ either one of the following laws:*

- **Excluded Middle.** $p \vee \neg p$
- **'Bad' Double Negation.** $\neg\neg p \to p.$

Proof. 21.1.2 already shows that the atomic instances of the Law of the Excluded Middle are enough to derive all tautologies.

21.1.3 seems instead to require the 'Bad' Double Negation Law as a schema, to step from provability of $\neg\neg\alpha$ to provability of α. But a simple induction on α in the language of $\neg$ and $\wedge$, which is classically adequate, shows that the full schema is implied by its atomic instances. For conjunctions, this follows from the fact that double negation completely filters through conjunctions (see the proof of 21.1.3). For negations, this follows from the following *Law of Triple Negation*:

$$\vdash_{\mathcal{N}} \neg\beta \leftrightarrow \neg\neg\neg\beta,$$

which is proved by the following:

$$\cfrac{\cfrac{\cfrac{[\gamma]^{(2)} \qquad [\neg\gamma]^{(1)}}{\bot}}{\neg\neg\gamma^{(1)}}}{\gamma^{(2)} \to \neg\neg\gamma.}$$

For $\gamma = \neg\beta$, this gives $\vdash_{\mathcal{N}} \neg\beta \to \neg\neg\neg\beta$. For $\gamma = \beta$, this gives $\vdash_{\mathcal{N}} \beta \to \neg\neg\beta$, and by contrapositive we then get $\vdash_{\mathcal{N}} \neg\neg\neg\beta \to \neg\beta$. □

Peirce's Law

The previous result showed that, from the point of view of the *complexity of formulas*, only the atomic instances of the Law of the Excluded Middle or of the 'Bad' Double Negation Law need to be assumed. We now take the complementary point of view of the *complexity of language*, and look for formulations of the two laws in terms of implication alone.

We first deal with the Law of the Excluded Middle, which is stated in the language of negation and disjunction. The latter can be reformulated in terms of implication, as follows:

- *negation*
  By definition, $\neg\alpha$ can be replaced by $\alpha \to \bot$.

- *disjunction*
  Since $\alpha \vee \gamma$ and $(\gamma \to \alpha) \to \alpha$ are classically equivalent (for example, because they have the same truth-tables), the former can be replaced by the latter.

Then we have the following classical equivalence:

$$(\alpha \vee \neg\alpha) \;\Leftrightarrow\; [(\alpha \to \bot) \to \alpha] \to \alpha.$$

The right-hand-side still involves the constant $\bot$, but is an instance of the following purely implicational schema, called **Peirce's Law**:

$$[(\alpha \to \beta) \to \alpha] \to \alpha.$$

There are two interesting facts about Peirce's Law. First, as a schema it implies every instance of the Law of the Excluded Middle. Second, each of its instances is classically provable, e.g. as follows:

$$\dfrac{\dfrac{\dfrac{\alpha \vdash_{\mathcal{SC}} \beta, \alpha}{\vdash_{\mathcal{SC}} \alpha \to \beta, \alpha} \qquad \alpha \vdash_{\mathcal{SC}} \alpha}{(\alpha \to \beta) \to \alpha \vdash_{\mathcal{SC}} \alpha}}{\vdash_{\mathcal{SC}} [(\alpha \to \beta) \to \alpha] \to \alpha.}$$

In other words, *Peirce's Law is an implicational translation of the Law of the Excluded Middle*, and it can be taken as an axiom schema in place of it. This is particularly useful if we are interested in the Classical Implicational Calculus, since then we get the following axiomation of it.

Theorem 21.2.4 Axiomatization of the Classical Implicational Calculus (Tarski and Bernays [19??]) *The Classical Implicational Calculus can be formulated as the system with the following axioms:*

1. $\gamma \to (\alpha \to \gamma)$
2. $[(\alpha \to (\gamma \to \delta)] \to [(\alpha \to \gamma) \to (\alpha \to \delta)]$
3. $[(\alpha \to \beta) \to \alpha] \to \alpha.$

Proof. We already know that the full Classical Propositional Calculus can be formulated by adding the Law of the Excluded Middle, and hence also Peirce's Law, to the axioms of the Intuitionistic Propositional Calculus. By the Normalization Theorem for the full Intuitionistic Propositional Calculus with axioms, any logical symbol or constant occurring in a normal proof must already occur in the conclusion or in the premises. Thus a valid classical implicational formula admits a proof in which the only logical symbol occurring is implication. Then such a proof is a proof in the Implicational Calculus from the axioms 1–3. □

Axiomatizations of the Classical Propositional Calculus

The trick just played with the Law of the Excluded Middle cannot be fully played with the 'Bad' Double Negation Law. Indeed, while we obviously have the following intuitionistic equivalence (by definition of $\neg$):

$$(\neg\neg\alpha \to \alpha) \; \Leftrightarrow \; [(\alpha \to \bot) \to \bot] \to \alpha,$$

the purely implicational schema

$$[(\alpha \to \beta) \to \beta] \to \alpha$$

is not classically valid, being false under any assignment that makes α false and β true. Thus the latter cannot be used for an axiomatization of the Classical Implicational Calculus, but the former can be used for a neat formulation of the full Classical Propositional Calculus.

Proposition 21.2.5 Axiomatization of the Full Classical Propositional Calculus. *The Classical Propositional Calculus can be formulated as the system with only one connective $\to$ and one constant $\bot$, and the following axioms:*

1. $\gamma \to (\alpha \to \gamma)$
2. $[(\alpha \to (\gamma \to \delta)] \to [(\alpha \to \gamma) \to (\alpha \to \delta)]$
3. $[(\alpha \to \bot) \to \bot] \to \alpha.$

Proof. By the adequacy of $\to$ and $\neg$. □

Obviously, various other formulations of the full Classical Propositional Calculus can be obtained by adding any of the Excluded Middle, 'Bad' Double Negation, 'Bad' Contrapositive, or Peirce's Laws to either the Natural Deduction or the Hilbert style axiomatizations of the Intuitionistic Propositional Calculus.

Any one of these additions gives *classical versions $\mathcal{NC}$ of Natural Deduction, and $\mathcal{HC}$ of Hilbert system*, for which we can prove the following result.

Proposition 21.2.6 Equivalence of Classical Systems. *For every Γ and α:*

$$\Gamma \vdash_{\mathcal{NC}} \alpha \; \Leftrightarrow \; \Gamma \vdash_{\mathcal{HC}} \alpha \; \Leftrightarrow \; \Gamma \vdash_{\mathcal{SC}} \alpha \; \Leftrightarrow \; \Gamma \models \alpha.$$

Proof. Since the same axiom was added to $\mathcal{N}$ and $\mathcal{H}$ to get the corresponding classical systems, the first equivalence follows from 1.2.2, which was proved for any set of assumptions.

In the second equivalence, one direction follows from the equivalence of $\mathcal{H}$ and $\mathcal{S}$ and any proof of the additional axiom $\mathcal{SC}$ (one is given in 21.2.4). The other direction follows from the last equivalence and 21.2.4.

The last equivalence was proved in the Classical Soundness and Completeness Theorem 19.1.4. □

Finally, a different formulation of $\mathcal{NC}$ can be obtained by adding the 'Bad' Double Negation Law to $\mathcal{N}$ not as an *axiom*, but as the following *rule*:

$$\frac{\Gamma \vdash_{\mathcal{NC}} \neg\neg\alpha}{\Gamma \vdash_{\mathcal{NC}} \alpha,}$$

corresponding to the possibility of extending proofs as follows:

$$\frac{\neg\neg\alpha}{\alpha.}$$

Then $\vdash_{\mathcal{NC}} \neg\neg\alpha \rightarrow \alpha$ becomes provable by →-introduction.

21.3 Extensions of the Classical Propositional Calculus

The results of this section provide a view of the relationships between the Classical and Intuitionistic Propositional Calculi opposite and complementary to the one exposed in the previous section.

Intuitionistic Propositional Calculus

For our present purposes, the basic observation is the following.

Theorem 21.3.1 (Gödel [1933]) *If α is a propositional formula in the language of $\neg$ and $\wedge$ alone, then α is classically valid if and only if it is intuitionistically valid, i.e.*

$$\models \alpha \;\Leftrightarrow \vdash_{\mathcal{N}} \alpha.$$

Proof. We only have to prove the left to right direction, since any formula provable in $\mathcal{N}$ is classically valid. We proceed inductively on the construction of α:

1. *letters*
   There is nothing to prove, since single letters are neither classically provable (because not valid), nor intuitionistically provable.

2. *negations*
   If $\alpha = \neg\beta$, then $\vdash_{\mathcal{N}} \neg\neg\neg\beta$ by 21.1.3, and $\vdash_{\mathcal{N}} \neg\beta$ by the Law of Triple Negation (proved in the proof of 21.2.3).

3. *conjunctions*
   If $\alpha = \beta \wedge \gamma$ and α is classically valid, then so are both β and γ. By the induction hypothesis, $\vdash_{\mathcal{N}} \beta$ and $\vdash_{\mathcal{N}} \gamma$, and by $\wedge$-introduction $\vdash_{\mathcal{N}} \beta \wedge \gamma$. □

The result is the best possible, in the sense that it fails when $\rightarrow$ or $\vee$ are added to the language. Peirce's Law is an example of a formula in the language of $\rightarrow$ that holds classically but not intuitionistically, and the Law of the Excluded Middle is an example in the language of $\neg$ and $\vee$.

Although the result fails for the full language of the Classical Propositional Calculus, this is not a restriction from a classical point of view, where every formula is equivalent to one in the language of $\neg$ and $\wedge$ alone. Thus the Intuitionistic Propositional Calculus can be seen as an *extension* of the classical one, in which the new connectives $\vee$ and $\rightarrow$ are *introduced* from scratch by new rules, instead of being *defined* from the old ones.

In a precise sense the Classical Propositional Calculus is thus *naturally* a subsystem of the intuitionistic one, since its connectives are interpreted truth-functionally from the very beginning (and the completeness proof of $\neg$, $\wedge$ is quite simple). On the other hand, the Intuitionistic Propositional Calculus is a subsystem of the classical one *unnaturally*, only if the meaning of its connectives (other than $\neg$ and $\wedge$) is brutally changed.

If we look at the full language, there is a class of formulas that are classically provable if and only if they are intuitionistically provable. They are the so-called *negative formulas*, i.e. those equivalent to ones that begin by a negation.

Proposition 21.3.2 (Glivenko [1929]) *If α is a propositional formula in the full language of $\neg$, $\wedge$, $\vee$ and $\rightarrow$, then its negation is classically valid if and only if it is intuitionistically valid, i.e.*

$$\models \neg\alpha \Leftrightarrow \vdash_{\mathcal{N}} \neg\alpha.$$

Proof. By 21.1.3,

$$\models \neg\alpha \Leftrightarrow \vdash_{\mathcal{N}} \neg\neg\neg\alpha.$$

And by the Triple Negation Law,

$$\vdash_{\mathcal{N}} \neg\neg\neg\alpha \Leftrightarrow \vdash_{\mathcal{N}} \neg\alpha. \quad \square$$

It follows in particular that a formula α is classically decidable, i.e. one of α and $\neg\alpha$ is provable, if and only if α is intuitionistically weakly decidable, i.e. one of $\neg\alpha$ and $\neg\neg\alpha$ is provable.

Minimal Propositional Calculus ⋆

Gödel's result 21.3.1 can be improved to show that the Classical Propositional Calculus is a subsystem not only of the intuitionistic one, but even of the weaker

Minimal Propositional Calculus, which is obtained from the intuitionistic one by retaining the constant $\bot$, but dropping the $\bot$-elimination rule.

Thus in the Minimal Propositional Calculus we do have negation, since $\neg\alpha$ can still be defined as $\alpha \to \bot$, but all its properties follow from properties of implication alone, because no special meaning is attached to the constant $\bot$.

This does not sound very promising, and we may think that only trivial properties of negation are provable in the Minimal Propositional Calculus. Contrary to this first impression, it turns out that most of the properties of negation are indeed minimal. For example, the following are, as the reader can easily verify:

1. **No Contradiction.** $\neg(\alpha \wedge \neg\alpha)$
2. **'Good' Contrapositive.** $(\alpha \to \beta) \to (\neg\beta \to \neg\alpha)$
3. **'Good' Double Negation.** $\alpha \to \neg\neg\alpha$
4. **Triple Negation.** $\neg\alpha \leftrightarrow \neg\neg\neg\alpha$
5. **'Good' De Morgan.** $\neg\alpha \wedge \neg\beta \leftrightarrow \neg(\alpha \vee \beta)$.

Moreover, the following:

6. **Weak Form of $\bot$-Elimination.** $\bot \to \neg\alpha$

also turns out to be minimally provable, being just an instance of $\beta \to (\alpha \to \beta)$, for $\beta = \bot$.

The following result is a strengthening of 21.3.1, obtained at no additional cost.

Corollary 21.3.3 *If α is a propositional formula in the language of $\neg$ and $\wedge$ alone, then α is classically valid if and only if it is minimally valid.*

Proof. By taking a close look at the proofs of 21.1.1, we see that we only made use of $\bot$-elimination in the proof of the following formation rule:

$$\neg\alpha \vdash_{\mathcal{N}} \alpha \to \beta.$$

In particular, the formation rules for $\neg$, $\wedge$ and $\vee$ have all been proved in the Minimal Propositional Calculus.

Since there was no use of $\bot$-elimination in the proof of 21.1.3, it follows that if α is a propositional tautology in the language of $\neg$, $\wedge$ and $\vee$, then $\neg\neg\alpha$ is provable in the Minimal Propositional Calculus.

Finally, since there was no use of $\bot$-elimination in the proof of 21.3.1, it follows that if α is a propositional tautology in the language of $\neg$ and $\wedge$, then α itself is provable in the Minimal Propositional Calculus. □

The Minimal Propositional Calculus has the following interesting property.

Proposition 21.3.4 Collapse of the Double Negation Translation (Friedman [1978]) *If α is a propositional formula not containing negations, then $\neg\neg\alpha$ is minimally provable if and only if so is α itself.*

Proof. By definition of negation,

$$\neg\neg\alpha = (\alpha \to \bot) \to \bot.$$

Since $\bot$ has no special meaning in the Minimal Propositional Calculus, by substituting any formula β for $\bot$ in the proof of $\neg\neg\alpha$ we still get a proof. If α contains no negation, then the substitution has no effect on it, and hence it produces a proof of

$$(\alpha \to \beta) \to \beta.$$

In particular, by letting $\beta = \alpha$ we get a proof of

$$(\alpha \to \alpha) \to \alpha.$$

Since the premise $\alpha \to \alpha$ is provable, so is the conclusion α. □

The restriction on α cannot be dropped. For example, $\neg\neg(p \vee \neg p)$ is provable in the Minimal Propositional Calculus by 21.1.3, but $p \vee \neg p$ is not. Actually, the latter is not even provable in the Intuitionistic Propositional Calculus.

The restriction on the Minimal Propositional Calculus cannot be dropped either. For example, the double negation of Peirce's Law is intuitionistically provable by 21.1.3, but Peirce's Law itself is not. Notice that the minimal analogue of 21.1.3 cannot be applied here to show that the double negation of Peirce's Law is minimally provable, because it only holds for formulas not containing $\to$.

21.4 Translations

In the present section we present a number of translations of the Classical into the Intuitionistic and Minimal Calculi, each with its own special properties.

Gödel's Translation

Gödel's result 21.3.1 deals with formulas in the restricted language of $\neg$ and $\wedge$ alone, but it can be reformulated as a result for formulas in the full language, after a translation of $\to$ and $\vee$ in terms of $\neg$ and $\wedge$.

Proposition 21.4.1 Gödel's Translation (Gödel [1933]) *If*

$$\begin{array}{rcl} p^* & = & p \\ (\neg\alpha)^* & = & \neg\alpha^* \\ (\alpha \wedge \beta)^* & = & \alpha^* \wedge \beta^* \\ (\alpha \vee \beta)^* & = & \neg(\neg\alpha^* \wedge \neg\beta^*) \\ (\alpha \rightarrow \beta)^* & = & \neg(\alpha^* \wedge \neg\beta^*), \end{array}$$

then, for every α,

$$\models \alpha \ \Leftrightarrow \vdash_{\mathcal{N}} \alpha^*.$$

Proof. The formula α^* is obtained by translating $\rightarrow$ and $\vee$ in the language of $\neg$ and $\wedge$ in a canonical way, and thus

$$\models \alpha \ \Leftrightarrow \models \alpha^* \ \Leftrightarrow \vdash_{\mathcal{N}} \alpha^*$$

because the translation is classically valid, and α^* is a formula in the language of $\neg$ and $\wedge$ alone. □

Gödel's Translation, while preserving *theorems*, does not preserve *consequences*. For example, $\neg\neg p \models p$, but $(\neg\neg p)^* \not\vdash_{\mathcal{N}} p^*$ because $\neg\neg p \not\vdash_{\mathcal{N}} p$, otherwise the 'Bad' Double Negation Law would be intuitionistically provable.

Gentzen's Translation

The weak point noted at the end of the previous subsection is repaired in the next result, which modifies Gödel's Translation by leaving implication unchanged, and doubly negating all propositional letters.

Proposition 21.4.2 Gentzen's Translation (Gentzen [1933]) *If*

$$\begin{array}{rcl} p^\circ & = & \neg\neg p \\ (\neg\alpha)^\circ & = & \neg\alpha^\circ \\ (\alpha \wedge \beta)^\circ & = & \alpha^\circ \wedge \beta^\circ \\ (\alpha \vee \beta)^\circ & = & \neg(\neg\alpha^\circ \wedge \neg\beta^\circ) \\ (\alpha \rightarrow \beta)^\circ & = & \alpha^\circ \rightarrow \beta^\circ, \end{array}$$

then, for every Γ *and* α,

$$\Gamma \models \alpha \ \Leftrightarrow \ \Gamma^\circ \vdash_{\mathcal{N}} \alpha^\circ,$$

where $\Gamma^\circ = \{\gamma^\circ : \gamma \in \Gamma\}$.

Proof. The result follows from the Double Negation Translation 21.1.5, by noticing that if β is a formula in the language of $\neg$, $\wedge$ and $\rightarrow$, then

$$\vdash_{\mathcal{N}} \beta^\circ \leftrightarrow \neg\neg\beta.$$

This is easily proved by induction on β. For letters, it holds by definition. For negations, il holds trivially. For conjunctions and implications, it holds because double negation filters through conjunction and implication. □

Notice that Gentzen's Translation of disjunction could also have been defined as

$$(\alpha \vee \beta)^\circ = \neg\neg(\alpha^\circ \vee \beta^\circ),$$

using the fact that the following 'Good' De Morgan's Law holds intuitionistically:

$$\neg\alpha \wedge \neg\beta \leftrightarrow \neg(\alpha \vee \beta).$$

Exercises 21.4.3 De Morgan's Laws.

a) $(\neg\alpha \wedge \neg\beta) \leftrightarrow \neg(\alpha \vee \beta)$ *holds intuitionistically (actually, minimally).*

b) $(\neg\alpha \vee \neg\beta) \rightarrow \neg(\alpha \wedge \beta)$ *holds intuitionistically (actually, minimally).*

c) $\neg(\alpha \wedge \beta) \rightarrow (\neg\alpha \vee \neg\beta)$ *fails intuitionistically.* (Hint: consider a Kripke model with three nodes $\emptyset$, 0 and 1, and such that $\mathcal{A}_\emptyset = \emptyset$, $\mathcal{A}_0 = \{\alpha\}$ and $\mathcal{A}_1 = \{\beta\}$.)

We can ask whether Gentzen's Translation of disjunction could simply be dropped, by defining

$$(\alpha \vee \beta)^\circ = \alpha^\circ \vee \beta^\circ.$$

This would make the translation a very simple homomorphism of propositional formulas, consisting only of the insertion of double negations in front of propositional letters, and leaving all connectives unchanged.

For the above proof to go through we would need *filtration of double negation through disjunction*, i.e.

$$\vdash_{\mathcal{N}} \neg\neg(\alpha \vee \beta) \leftrightarrow (\neg\neg\alpha \vee \neg\neg\beta).$$

But while the right to left direction holds, as it can easily be checked, the left to right direction fails intuitionistically (see 21.4.5).

Actually, it is not only the particular proof considered above that fails, but the result itself. Otherwise, from $\models \alpha \vee \neg\alpha$ we would get $\vdash_{\mathcal{N}} \neg\neg\alpha \vee \neg\neg\neg\alpha$, and thus $\vdash_{\mathcal{N}} \neg\neg\alpha \vee \neg\alpha$ by the Triple Negation Law. But this *weak form of the Law of the Excluded Middle* fails intuitionistically (see 21.4.5).

Gentzen's Translation, while preserving *consequences*, does not preserve *substitution*: for example, $p^* = \neg\neg p$ but $(\alpha \wedge \beta)^* \neq \neg\neg(\alpha^* \wedge \beta^*)$.

Kolmogorov's Translation

The weak point noted at the end of the previous subsection is repaired in the next result, which modifies Gentzen's Translation by doubly negating all propositional letters and subformulas (except for negation, whose double negation has no effect by the Triple Negation Law).

Proposition 21.4.4 Kolmogorov's Translation (Kolmogorov [1925]) *If*

$$\begin{array}{rcl} p^\diamond & = & \neg\neg p \\ (\neg\alpha)^\diamond & = & \neg\alpha^\diamond \\ (\alpha\wedge\beta)^\diamond & = & \neg\neg(\alpha^\diamond\wedge\beta^\diamond) \\ (\alpha\vee\beta)^\diamond & = & \neg\neg(\alpha^\diamond\vee\beta^\diamond) \\ (\alpha\rightarrow\beta)^\diamond & = & \neg\neg(\alpha^\diamond\rightarrow\beta^\diamond), \end{array}$$

then, for every Γ *and* α,

$$\Gamma\models\alpha \ \Leftrightarrow\ \Gamma^\diamond\vdash_{\mathcal{N}}\alpha^\diamond,$$

where $\Gamma^\diamond=\{\gamma^\diamond:\gamma\in\Gamma\}$.

Proof. Instead of indirectly relying on previous results such as 21.1.3 or 21.3.1, we directly show how to transform classical into intuitionistic proofs. This is done not only to provide an independent proof of the result, but also to derive stronger consequences.

First notice that

$$\vdash_{\mathcal{N}}\alpha^\diamond\leftrightarrow\neg\neg\alpha^\diamond,$$

since $\alpha^\diamond$ always begins with $\neg\neg$, and by the triple negation law four negations are equivalent to two.

We now start from a proof of $\Gamma\vdash\alpha$ in the system $\mathcal{NC}$, and produce a proof of $\Gamma^\diamond\vdash\alpha^\diamond$ in the system $\mathcal{N}$ by substituting each formula in the given proof by its translation, and inserting proofs of appropriate lemmas when needed. There are three sets of cases:

- *axioms*
  Since $(p\vee\neg\neg p)^\diamond=\neg\neg(\neg\neg p\vee\neg p)$ by the Triple Negation Law, the translations of the axioms are provable, by a proof similar to the of $\neg\neg(\neg\neg p\vee\neg p)$ given in the proof of 21.1.3.

- *introduction rules*
  A $\wedge$-introduction step

  $$\frac{\alpha \qquad \beta}{\alpha\wedge\beta}$$

  becomes

  $$\begin{array}{c} \alpha^\diamond \qquad \beta^\diamond \\ \hline \alpha^\diamond\wedge\beta^\diamond \\ \hline\hline \neg\neg(\alpha^\diamond\wedge\beta^\diamond) \\ \hline (\alpha\wedge\beta)^\diamond \end{array}$$

  by the induction hypothesis, an instance of a proof of $\gamma\rightarrow\neg\neg\gamma$, and definition of $^\diamond$. Here and in the following, we indicate by a double line the fact that we step from the top to the bottom not directly, but rather by inserting a proof.

A $\vee$-introduction step

$$\begin{array}{c} \alpha \\ \hline \alpha \vee \beta \end{array}$$

becomes

$$\begin{array}{c} \alpha^{\diamond} \\ \hline \alpha^{\diamond} \vee \beta^{\diamond} \\ \hline\hline \neg\neg(\alpha^{\diamond} \vee \beta^{\diamond}) \\ \hline (\alpha \vee \beta)^{\diamond} \end{array}$$

by the induction hypothesis, an instance of a proof of $\gamma \to \neg\neg\gamma$, and definition of $^{\diamond}$.

A $\to$-introduction step

$$\begin{array}{c} [\alpha] \\ \mathcal{D} \\ \beta \\ \hline \alpha \to \beta \end{array}$$

becomes

$$\begin{array}{c} [\alpha^{\diamond}] \\ \mathcal{D}^{\diamond} \\ \beta^{\diamond} \\ \hline \alpha^{\diamond} \to \beta^{\diamond} \\ \hline\hline \neg\neg(\alpha^{\diamond} \to \beta^{\diamond}) \\ \hline (\alpha \to \beta)^{\diamond} \end{array}$$

by the induction hypothesis, an instance of a proof of $\gamma \to \neg\neg\gamma$, and definition of $^{\diamond}$.

- *elimination rules*
  A $\wedge$-elimination step

$$\begin{array}{c} \alpha \wedge \beta \\ \hline \alpha \end{array}$$

becomes

$$\begin{array}{c} \begin{array}{c} \begin{array}{c} [\alpha^{\diamond} \wedge \beta^{\diamond}]^{(1)} \\ \hline \alpha^{\diamond} \end{array} \quad [\neg\alpha^{\diamond}]^{(2)} \\ \hline \bot \\ \hline \neg(\alpha^{\diamond} \wedge \beta^{\diamond})^{(1)} \end{array} \qquad \begin{array}{c} (\alpha \wedge \beta)^{\diamond} \\ \hline\hline \neg\neg(\alpha^{\diamond} \wedge \beta^{\diamond}) \end{array} \\ \hline \bot \\ \hline (\neg\neg\alpha^{\diamond})^{(2)} \\ \hline\hline \alpha^{\diamond} \end{array}$$

by definition of $^{\diamond}$, and the equivalence between $\neg\neg\alpha^{\diamond}$ and $\alpha^{\diamond}$.

A $\vee$-elimination step

$$\dfrac{\alpha\vee\beta \qquad \begin{array}{c}[\alpha]\\ \mathcal{D}_1\\ \gamma\end{array} \qquad \begin{array}{c}[\beta]\\ \mathcal{D}_2\\ \gamma\end{array}}{\gamma}$$

becomes

$$\dfrac{\underline{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{[\alpha^\diamond\vee\beta^\diamond]^{(1)} \qquad \begin{array}{c}[\alpha^\diamond]\\ \mathcal{D}_1^\diamond\\ \gamma^\diamond\end{array} \qquad \begin{array}{c}[\beta^\diamond]\\ \mathcal{D}_2^\diamond\\ \gamma^\diamond\end{array}}{\gamma^\diamond} \qquad [\neg\gamma^\diamond]^{(2)}}{\bot}}{\neg(\alpha^\diamond\vee\beta^\diamond)^{(1)}} \qquad \dfrac{\underline{(\alpha\vee\beta)^\diamond}}{\neg\neg(\alpha^\diamond\vee\beta^\diamond)}}{\bot}}{(\neg\neg\gamma^\diamond)^{(2)}}}}{\gamma^\diamond}$$

by definition of $^\diamond$, and the equivalence between $\neg\neg\gamma^\diamond$ and $\gamma^\diamond$.

A $\rightarrow$-elimination step

$$\dfrac{\alpha \qquad \alpha\rightarrow\beta}{\beta}$$

becomes

$$\dfrac{\underline{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\alpha^\diamond \qquad [\alpha^\diamond\rightarrow\beta^\diamond]^{(1)}}{\beta^\diamond} \qquad [\neg\beta^\diamond]^{(2)}}{\bot}}{\neg(\alpha^\diamond\rightarrow\beta^\diamond)^{(1)}} \qquad \dfrac{\underline{(\alpha\rightarrow\beta)^\diamond}}{\neg\neg(\alpha^\diamond\rightarrow\beta^\diamond)}}{\bot}}{(\neg\neg\beta^\diamond)^{(2)}}}}{\beta^\diamond}$$

by definition of $^\diamond$, and the equivalence between $\neg\neg\beta^\diamond$ and $\beta^\diamond$.

Finally, we deal with a $\bot$-elimination step

$$\dfrac{\bot}{\alpha.}$$

First,

$$\dfrac{\underline{\dfrac{\dfrac{\bot}{\neg\alpha^\diamond\rightarrow\bot}}{\neg\neg\alpha^\diamond}}}{\alpha^\diamond}$$

by →-introduction, definition of $\neg$ and the equivalence between $\neg\neg\alpha^{\diamond}$ and $\alpha^{\diamond}$. Then, by two applications of the 'good' contrapositive law,

$$\begin{array}{c}\neg\neg\bot\\ \hline\hline \neg\neg\alpha^{\diamond},\end{array}$$

and hence

$$\begin{array}{c}\bot^{\diamond}\\ \hline\hline \alpha^{\diamond}\end{array}$$

by the equivalence between $\neg\neg\alpha^{\diamond}$ and $\alpha^{\diamond}$. □

Notice how the lemmas we need to insert in the tranformed proof are quite trivial and uniform. They use introductions and instances of $\gamma \rightarrow \neg\neg\gamma$ in one direction, eliminations and instances of $\neg\neg\alpha^{\diamond} \leftrightarrow \alpha^{\diamond}$ in the other direction. Moreover, they never use the $\bot$-elimination rule, even in the transformation of a $\bot$-elimination step, and are thus proofs in the Minimal Propositional Calculus.

We could easily recast the proofs of the Double Negation or Gentzen's Translations in the same framework of proof transformation just used for Kolmogorov's Translation. Then we would need to insert proofs of filtration of $\neg\neg$ through connectives. However, some of these proofs (given in the proofs of 21.1.3 and 21.1.5) are not trivial, one (filtration through →) uses $\bot$-elimination, and one (filtration through $\vee$) simply fails. In particular, on the one hand, the given transformation cannot deal with proofs in the full system, only in the system without disjunction. On the other hand, it only provides a translation into the Intuitionistic, not into the Minimal Propositional Calculus.

All things said, Kolmogorov's Translation is thus the most efficient among the ones we considered. From the point of view of *formula transformation*, it preserves substitution. From the point of view of *proof transformation*, it restrains the size of the proof by inserting only very simple and uniform lemmas, and it translates into the Minimal, not just into the Intuitionistic Propositional Calculus.

But even Kolmogorov's Translation is not the last word. On the one hand, it uses the maximum possible number of $\neg\neg$, namely, the number of occurrences of letters and connectives in a given formula. On the other hand, it does not provide an *isomorphism of proofs* of the formula in $\mathcal{NC}$, and of the translation in $\mathcal{N}$. Girard [199?] has devised a translation that minimizes the number of $\neg\neg$, and does provide such an isomorphism.

Weak Excluded Middle ⋆

The next result shows that the two principles whose failure we mentioned in the discussion of Gentzen's Translation are actually intuitionistically equivalent, and that the addition of any of them to the Intuitionistic Propositional Calculus provides an *Intermediate Propositional Calculus* stronger than the intuitionistic, but weaker than the classical one.

Proposition 21.4.5 *The following two schemata:*

1. **Weak Excluded Middle.** $\neg\alpha \vee \neg\neg\alpha$
2. **Filtration of $\neg\neg$ through $\vee$.** $\neg\neg(\alpha \vee \beta) \to (\neg\neg\alpha \vee \neg\neg\beta)$

are intuitionistically equivalent, not valid, and strictly weaker than the Law of Excluded Middle.

Proof. 1 implies 2 as follows. By 1, both $\neg\alpha \vee \neg\neg\alpha$ and $\neg\beta \vee \neg\neg\beta$ hold. Hence so does the following disjunction, by the Distributivity Laws of $\wedge$ and $\vee$ (which are intuitionistically valid):

$$(\neg\alpha \wedge \neg\beta) \vee (\neg\alpha \wedge \neg\neg\beta) \vee (\neg\neg\alpha \wedge \neg\beta) \vee (\neg\neg\alpha \wedge \neg\neg\beta).$$

To prove 2, assume $\neg\neg(\alpha \vee \beta)$. By the 'Good' De Morgan's Law, $\neg(\neg\alpha \wedge \neg\beta)$ holds, and hence the first disjunct above fails. But each of the three remaining ones implies $\neg\neg\alpha \vee \neg\neg\beta$, because it implies one of $\neg\neg\alpha$ and $\neg\neg\beta$.

2 implies 1 because $\neg\neg(\alpha \vee \neg\alpha)$ holds by 21.1.3, and by 2 (with $\beta = \neg\alpha$) we get $\neg\neg\alpha \vee \neg\neg\neg\alpha$, from which 1 follows by the Triple Negation Law.

Because of the equivalences just proved, we can concentrate on the Law of the Weak Excluded Middle and notice that:

- It is implied by the Law of the Excluded Middle, being a special case of it for negated formulas.
- It does not imply the Law of the Excluded Middle, because the former holds but the latter fails in a Kripke model with two nodes $\emptyset$ and 0, and such that $\mathcal{A}_\emptyset = \emptyset$ and $\mathcal{A}_0 = \{\alpha\}$.
- It is not intuitionistically valid, because it does not hold in a Kripke model with three nodes $\emptyset$, 0 and 1, and such that $\mathcal{A}_\emptyset = \mathcal{A}_0 = \emptyset$, and $\mathcal{A}_1 = \{\alpha\}$. □

Actually, a number of intuitionistic principles turn out to be equivalent to the two just discussed. The next exercise provides an additional one, and Johnstone [1979] discusses the matter in detail.

Exercise 21.4.6 The 'Bad' De Morgan law. We have seen that the 'Good' De Morgan Law is intuitionistically valid. We now consider the 'bad' one, namely

$$\neg(\alpha \wedge \beta) \to \neg\alpha \vee \neg\beta.$$

The 'Bad' De Morgan's Law is intuitionistically equivalent to the Weak Law of the Excluded Middle. (Hint: it is easier to prove the equivalence with the filtration of $\neg\neg$ through $\vee$. In one direction, we proceed as follows:

$$\begin{array}{lll} & \neg(\alpha \wedge \beta) & \\ \Rightarrow & \neg\neg\neg(\alpha \wedge \beta) & \text{by the Triple Negation Law} \\ \Rightarrow & \neg(\neg\neg\alpha \wedge \neg\neg\beta) & \text{by filtration of } \neg\neg \text{ through } \wedge \\ \Rightarrow & \neg\neg(\neg\alpha \vee \neg\beta) & \text{by the 'Good' De Morgan's Law} \\ \Rightarrow & \neg\neg\neg\alpha \vee \neg\neg\neg\beta & \text{by filtration of } \neg\neg \text{ through } \vee \\ \Rightarrow & \neg\alpha \vee \neg\beta & \text{by the Triple Negation Law.} \end{array}$$

In the other direction, from $\neg\neg(\alpha \vee \beta)$ we get $\neg(\neg\alpha \wedge \neg\beta)$ by the 'Good' De Morgan's Law, and then $\neg\neg\alpha \vee \neg\neg\beta$ by the 'Bad' De Morgan's Law.)

Exercise 21.4.7 *A Kripke model forces all instances of the Weak Law of the Excluded Middle if and only if it (is equivalent to one that) has a greatest element.* (Hint: suppose the model $\mathcal{A}$ has a greatest element a. By definition of forcing a negation, if $\emptyset$ does not force $\neg\alpha$, then there is a state forcing α. By monotonicity of forcing, a must force α too, and then $\emptyset$ forces $\neg\neg\alpha$.

Conversely, suppose a greatest element cannot be added. There must be two elements a and b and a letter p, such that p is forced above a but not above b. Then neither $\neg p$ nor $\neg\neg p$ can be forced, the former because p is forced above a, and the latter because p is not forced above b.)

Exercises 21.4.8 Dummett Intermediate Logic (Skolem [1913], Gödel [1932], Dummett [1959], Horn [1962])

a) *The following two schemata:*

1. $(\alpha \to \beta) \vee (\beta \to \alpha)$

2. $(\alpha \wedge \beta \to \gamma) \to [(\alpha \to \gamma) \vee (\beta \to \gamma)]$

are intuitionistically equivalent. (Hint: 1 implies 2 by cases, because if $\alpha \to \beta$ then $\alpha \to \alpha \wedge \beta$, and from $\alpha \wedge \beta \to \gamma$ we get $\alpha \to \gamma$. Similarly, if $\beta \to \alpha$, then $\beta \to \gamma$.

2 implies 1 by letting γ be $\alpha \wedge \beta$. Then $(\alpha \to \alpha \wedge \beta) \vee (\beta \to \alpha \wedge \beta)$, from which 1 follows by $\alpha \wedge \beta \to \alpha$ and $\alpha \wedge \beta \to \beta$.)

b) *The previous two schemata are intuitionistically implied by the Law of the Excluded Middle, but not conversely.* (Hint: since $\alpha \to (\beta \to \alpha)$ and $\neg\alpha \to (\alpha \to \beta)$ hold intuitionistically, 1 follows from $\alpha \vee \neg\alpha$ by cases.

The Kripke model with two nodes $\emptyset$ and 0, and such that $\mathcal{A}_\emptyset = \emptyset$ and $\mathcal{A}_0 = \{\alpha\}$, shows that the converse implication fails.)

c) *The previous two schemata intuitionistically imply the Weak Law of the Excluded Middle, but not conversely.* (Hint: the implication follows from 21.4.6, since 2 reduces to the 'Bad' De Morgan Law when $\gamma = \perp$.

The Kripke model with four nodes 0, a, b and 1, with a and b incomparable, and such that $\mathcal{A}_0 = \emptyset$, $\mathcal{A}_a = \{\alpha\}$, $\mathcal{A}_b = \{\beta\}$ and $\mathcal{A}_1 = \{\alpha, \beta\}$, shows that the converse implication fails.)

d) *A Kripke model (equivalent to one) with a smallest element forces all instances of the previous two schemata if and only if it is linear.* (Hint: by 5.3.2.d and the definition of forcing for disjunction. The Kripke model with only two incomparable elements a and b, and such that $\mathcal{A}_a = \{\alpha\}$ and $\mathcal{A}_b = \{\beta\}$, shows that the condition on the smallest element is necessary.)

e) *A formula is an intuitionistic consequence of the previous two schemata if and only if it is forced in every linear Kripke model.* (Hint: by part d), since the models provided by the Intuitionistic Completeness Theorem all have a smallest element.)

f) *A formula is an intuitionistic consequence of the previous two schemata if and only if it holds in every linear Heyting algebra.* (Hint: by part e) and 5.3.2.)

Intuitionistic Propositional Logic can thus be extended to at least three, increasingly comprehensive logics, by respectively adding to its axioms $\neg\alpha \vee \neg\neg\alpha$, $(\alpha \to \beta) \vee (\beta \to \alpha)$, or $\alpha \vee \neg\alpha$. In the first case one gets the intermediate logic studied in this section, in the second Dummett Logic, and in the last full Classical Propositional Logic.

21.5 Classical and Intuitionistic Connectives

We have considered in this chapter a number of classically valid equivalences that are intuitionistically only 'half valid', in the sense that one implication is valid but the other one is not.

This is so, in particular, for the usual classical reductions of all connectives (and, actually, of all truth-valued functions) to negation and either conjunction or disjunction, as the next result shows.

Proposition 21.5.1 *The following formulas hold intuitionistically (actually, in the first two cases, minimally):*

$$(\alpha \wedge \beta) \quad \to \quad \neg(\neg\alpha \vee \neg\beta) \tag{21.1}$$

$$(\alpha \vee \beta) \quad \to \quad \neg(\neg\alpha \wedge \neg\beta) \tag{21.2}$$

$$(\alpha \to \beta) \quad \leftarrow \quad (\neg\alpha \vee \beta). \tag{21.3}$$

Proof. To prove 1:

$$\dfrac{\dfrac{[\neg\alpha \vee \neg\beta]^{(2)} \qquad \dfrac{[\neg\alpha]^{(1)} \quad \dfrac{\alpha \wedge \beta}{\alpha}}{\bot} \qquad \dfrac{[\neg\beta]^{(1)} \quad \dfrac{\alpha \wedge \beta}{\beta}}{\bot}}{\bot^{(1)}}}{\neg(\neg\alpha \vee \neg\beta)^{(2)}.}$$

To prove 2:

$$\dfrac{\alpha\vee\beta \qquad \dfrac{[\alpha]^{(1)} \qquad \dfrac{[\neg\alpha\wedge\neg\beta]^{(2)}}{\neg\alpha}}{\bot} \qquad \dfrac{[\beta]^{(1)} \qquad \dfrac{[\neg\alpha\wedge\neg\beta]^{(2)}}{\neg\beta}}{\bot}}{\dfrac{\bot^{(1)}}{\neg(\neg\alpha\wedge\neg\beta)^{(2)}.}}$$

To prove 3:

$$\dfrac{\dfrac{\neg\alpha\vee\beta \qquad \dfrac{\dfrac{[\neg\alpha]^{(1)} \quad [\alpha]^{(2)}}{\bot}}{\beta} \qquad [\beta]^{(1)}}{\beta^{(1)}}}{\alpha^{(2)}\to\beta.}$$

Since the step from $\bot$ to β makes an appeal to the $\bot$-Rule, the previous proof is not minimal. No such an appeal is made instead in the first two proofs, which are thus minimal. □

Obviously, if any of the previous implication were reversible then the relative connective would be the definable in terms of the others. Unfortunately, this is not the case.

Proposition 21.5.2 *None of $\neg$, $\wedge$, $\vee$ and $\to$ is intuitionistically definable from the remaining ones.*

Proof. Since $\neg$ is defined in terms of $\bot$ and $\to$, it is enough to prove that none of $\bot$, $\wedge$, $\vee$ and $\to$ is definable from the remaining ones.

- *falsity*
  Consider the Heyting algebra $\{0,1\}$, and the environment ρ which sends all propositional letters to 1. Then $[\![\bot]\!]_\rho = 0$, but $[\![\alpha]\!]_\rho = 1$ for any formula α involving only $\wedge$, $\vee$ and $\to$.

  Since $\{0,1\}$ is actually a Boolean algebra, this shows that $\bot$ is not even classically definable in terms of $\wedge$, $\vee$ and $\to$.

- *conjunction*
  Consider the Heyting algebra obtained as the product of $\{0,1\}$ and $\{0,a,1\}$, and the environment ρ which sends the letter p to $(0,1)$ and all the others to $(1,a)$. If q is different from p, then $[\![p\wedge q]\!]_\rho = (0,a)$, but $[\![\alpha]\!]_\rho \neq (0,a)$ for any formula α involving only p and q, as well as $\bot$, $\vee$ and $\to$.

- *disjunction*
  Consider the Heyting algebra

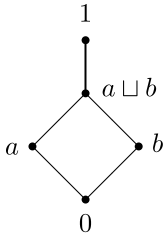

and any environment ρ which sends the letter p to a, and all the others to b. If q is different from p, then $\llbracket p \vee q \rrbracket_\rho = a \sqcup b$, but $\llbracket \alpha \rrbracket_\rho \neq a \sqcup b$ for any formula α involving only p and q, as well as $\perp$, $\wedge$ and $\rightarrow$ (notice, in particular, that $a \Rightarrow 0 = \sim a = b$).

- *implication*
  Consider the Heyting algebra obtained as the product of $\{0, a, 1\}$ and $\{0, a, 1\}$, and the environment ρ which sends the letter p to (a, a) and all the others to $(1, a)$. If q is different from p, then $\llbracket p \rightarrow q \rrbracket_\rho = (a, 1)$, but $\llbracket \alpha \rrbracket_\rho \neq (a, 1)$ for any formula α involving only p and q, as well as $\perp$, $\wedge$ and $\vee$. □

Thus in intuitionistic logic every connective plays its own individual role, and it is not reducible to the other ones as in classical logic. As one might expect, other indipendent intuitionistic connectives can be added to the four usual ones: it follows from Gödel [1932] that there are actually infinitely many possibilities, and from McKinsey and Tarsky [1946] that this is true even of the unary connectives, i.e. those built from a single propositional letter.

In terms of the Lindenbaum algebra this corresponds to the fact that, while in classical logic the *free Boolean algebra on one generator* has only four elements, namely:

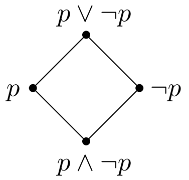

in intuitionistic logic the *free Heying algebra on one generator* has infinitely many elements, with the following structure characterized by Nishimura [1960]:

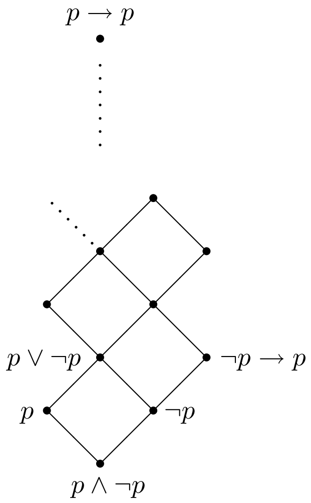
$p \to p$
$p \vee \neg p$
$\neg p \to p$
p
$\neg p$
$p \wedge \neg p$

Bibliography

Abian, S., and Brown, A.B.

[1961] A theorem on partially ordered sets with applications to fixed-point theorems, *Can. J. Math.* 13 (1961) 78–83.

Alexandrov, P.S.

[1937] Diskrete Räume, *Math. Sb.* 43 (1937) 501–519.

Baeten, J., and Boerboom, B.

[1979] Ω can be anything it shouldn't be, *Indag. Math.* 41 (1979) 111–120.

Barendregt, H.

[1981] *The Lambda Calculus*, North Holland, 1981 (Second edition, 1984).

Beth, E.W.

[1956] Semantic construction of intuitionistic logic, *Kon. Nederl. Akad. Wetensch.* 19 (1956) 357–388.

Birkhoff, G.

[1933] On the combination of subalgebras, *Proc. Cambr. Phil. Soc.* 29 (1933) 441–464.

[1940] *Lattice theory*, American Mathematical Society, 1940.

Birkhoff, G., and Frink, O.

[1948] Representation of lattices by sets, *Trans. Am. Math. Soc.* 64 (1948) 299–316.

Büchi, J.R.

[1952] Representation of complete lattices by sets, *Portug. Math.* 11 (1952) 151–167.

Cantor, G.

[1874] Über eine Eigenschaft des Inbegriffes aller reellen algebraischen Zahlen, *J. Math.* 77 (1874) 258–262.

Cartan, H.

[1937] Théorie des filtres, *Compt. Rend. Acad. Sci.* 205 (1937) 595–598.

[1937a] Filtres et ultrafilters, *Compt. Rend. Acad. Sci.* 205 (1937) 777–779.

Church, A.

[1933] A set of postulates for the foundation of logic (second paper), *Ann. Math.* 34 (1933) 839–864.

[1936] An unsolvable problem of elementary number theory, *Am. J. Math.* 58 (1936) 345–363.

[1936a] A note on the Eintscheidungsproblem, *J. Symb. Log.* 1 (1936) 40–41.

[1941] *The Calculi of Lambda Conversion*, Princeton University Press, 1941.

Church, A., and Rosser, B.J.

[1936] Some properties of conversion, *Trans. Am. Math. Soc.* 39 (1936)472–482.

Cohen, P.J.

[1963] The independence of the continuum hypothesis, *Proc. Nat. Acad. Sci.* 50 (1963) 1143–1148.

Collins, G.E.

[1954] Distributivity and an axiom of choice, *J. Symb. Log.* 19 (1954) 275–277.

Coppo, M., and Dezani, M.

[1980] An extension of basic functionality theory for λ-calculus, *Notre Dame J. Form. Log.* 21 (1980) 685–693.

Coppo, M., Dezani, M., and Venneri, B.

[1981] Functional characters of solvable terms, *Zeit. Math. Log. Grund. Math.* 27 (1981) 45–58.

Curry, H.B.

[1930] Grundlagen der Kombinatorischen Logik, *Am. J. Math.* 52 (1930) 789–834.

[1942] The inconsistency of certain formal logics, *J. Symb. Log.* 7 (1942) 115–117.

[1969] The undecidability of λK-conversion, in *Foundations of Mathematics*, Bulloff et al. eds., Springer, 1969, pp. 10–14.

Curry, H.B., and Feys, R.

[1958] *Combinatory logic*, North Holland, 1958.

Curry, H.B., Hindley, J.R., and Seldin, J.P.

[1972] *Combinatory logic*, volume II, North Holland, 1972.

Day, B.J., and Kelly, G.M.

[1970] On topological quotients maps preserved by pullbacks or products, *Proc. Cambr. Phil. Soc.* 67 (1970) 553–558.

Dragalin, A.G.

[1968] The computation of primitive recursive terms of finite type, and primitive recursive realization, *Zap. Nau. Sem. Lenin.* 8 (1968) 32–45.

Dummett, M.A.

[1959] A propositional calculus with denumerable matrix, *J. Symb. Log.* 24 (1959) 97–106.

Dummett, M.A., and Lemmon, E.J.

[1959] Modal logics between S4 and S5, *Zeit. Math. Log. Grund. Math.* 5 (1959) 250–264.

Engeler, E.

[1981] Algebras and combinators, *Alg. Univ.* 13 (1981) 389–392.

Fitting, M.

[1983] *Proof methods for modal and intuitionistic logics*, Reidel, 1983.

Freyd, P.J.

[1964] *Abelian categories*, Harper and Row, 1964.

Friedberg, R.M., and Rogers, H.

[1959] Reducibilities and completeness for sets of integers, *Zeit. Math. Log. Grund. Math.* 5 (1959) 117–125.

Friedman, H.

[1975] Equality between functionals, *Springer Lect. Not. Math.* 453 (1975) 22–37.

[1978] Classically and intuitionistically provably recursive functions, *Springer Lect. Not. Math.* 669 (1978) 21–28.

Gandy, R.O.

[1980] Proofs of strong normalization, in Seldin and Hindley [1980], pp. 457–477.

Gentzen, G.

[1935] Untersuchungen über das logische Schliessen, *Math. Zeit.* 39 (1935) 176–210, 405–431, transl. in [1969], pp. 68–131.

[1969] *Collected papers*, North Holland, 1969.

Gierz, G., Hofmann, K.H., Keimel, K., Lawson, J.D., Mislove, M., Scott, D.S.

[1980] *A compendium of continuous lattices*, Springer, 1980.

Gödel, K.

[1932] Zum intuitionistischen Aussagenkalkül, *Anz. Akad. Wiss. Wien* 69 (1932) 65–66, transl. in [1986], pp. 222–225.

[1986] *Collected works, volume I*, Oxford University Press, 1986.

Grassmann, H.

[1861] *Lehrbuch der Arithmetik für höhere Lehranstalten*, Berlin, 1861.

Grothendieck, A., and Dieudonné, J.

[1960] *Élements de géometrie algébrique*, Volume I, Springer Verlag, 1960.

Gunter, C.A., and Scott, D.S.

[1990] Semantic domains, in *Handbook of Theoretical Computer Science, volume B*, van Leeuwen ed., North Holland, 1990, pp. 635–674.

Halpern, J.D.

[1964] The independence of the axiom of choice from the Boolean prime ideal theorem, *Fund. Math.* 55 (1964) 57–66.

Henkin, L.

[1954] Boolean representation through propositional calculus, *Fund. Math.* 41 (1954) 89–96.

Herbrand, J.

[1928] Sur la théorie de la démonstration, *Compt. Rend. Acad. Sci.* 186 (1928) 1274–1276.

Heyting, A.

[1930] Die formalen Regeln der intuitionistichen Logik, *Sitzung. Preuss. Akad. Wiss. Phys. Math. Klasse*, 1930, pp. 42–71.

Hinata, S.

[1967] Calculability of primitive recursive functionals of finite type, *Sci. Rep. Tokyo Kyoi. Daig.* 9 (1967) 218–235.

Hinatani, Y.

[1966] Calculabilité des fonctionnels recursives primitives de type fini sur les nombres naturels, *Ann. Jap. Ass. Phil. Sci.* 3 (1966) 19–30.

Hindley, J.R.

[1997] *Basic Simple Type Theory*, Cambridge University Press, 1997.

Hindley, J.R., and Seldin, J.P.

[1986] *Introduction to Combinators and λ-Calculus*, Cambridge University Press, 1986.

Hofmann, K.H., and Keimel, K.

[1972] *A general character theory for partial ordered sets and lattices*, Memoirs of the American Mathematical Society n. 122, 1972.

Hofmann, K.H., and Lawson, J.D.

[1978] The spectral theory of distributive continuous lattices, *Trans. Amer. Math. Soc.* 246 (1978) 285–310.

Hofmann, R.E.

[1981] Continuous posets, prime spectra of completely distributive complete lattices, and Hausdorff compactifications, *Springer Lect. Not. Math.* 871 (1981) 159–208.

Horn, A.

[1962] Logic with truth-values in a linearly ordered Heyting algebra, *J. Symb. Log.* 27 (1962) 159–170.

Howard, W.

[1970] Assignments of ordinals to terms for primitive recursive functionals of finite type, in *Intuitionism and Proof Theory*, Kino et al. eds., North Holland, 1970, pp. 443–458.

Huet,G.

[1980] Confluent reductions, *J. Ass. Comp. Mach.* 27 (1980) 797–821.

Hughes, G.E., and Cresswell, M.J.

[1968] *An introduction to modal logic*, Methuen, 1968.

Huntington, E.V.

[1904] Sets of independent postulates for the algebra of logic, *Trans. Am. Math. Soc.* 5 (1904) 288–309.

Jaskowski, S.

[1936] Recherches sur le système de la logique intuitionniste, *Actes Congr. Intern. Phil. Scient.* 6 (1936) 58–61.

Johnstone, P.T.

[1979] Conditions related to De Morgan's law, *Springer Lect. Not. Math.* 753 (1979) 479–491.

[1982] *Stone spaces*, Cambridge University Press, 1982.

[1983] The point of pointless topology, *Bull. Am. Math. Soc.* (1983) 41–53.

Kamara, M.

[1978] Treillis continus et treillis complètement distributifs, *Semigr. Forum* 16 (1978) 387–388.

Kan, D.M.

[1958] Adjoint functors, *Trans. Am. Math. Soc.* 87 (1958) 294–329.

Kleene, S.K.

[1935] A theory of positive integers in formal logic, *Am. J. Math.* 57 (1935) 153–173, 219–244.

[1936] λ-definability and recursiveness, *Duke Math. J.* 2 (1936)340–353.

Klimovsky, G.

[1958] El teorema de Zorn y la existencia de filtros y ideales maximales en los reticulados distributivos, *Rev. Un. Mat. Argen.* 18 (1958) 160–164.

Klop, J.W.

[1980] *Combinatory reduction systems*, Amsterdam Mathematisch Centrum, 1980.

Knaster, B.

[1928] Un théorème sur les fonctions d'ensembles, *Ann. Soc. Polon. Math.* 6 (1928) 133–134.

Kreisel, G., and Putnam, H.

[1957] Eine Unableitbarkeitsbeweismethode für den intuitionistichen Aussagenkalkül, *Arch. Math. Log.* 3 (1957) 74–78.

Kripke, S.

[1963] Semantical considerations on modal and intuitionistic logic, *Acta Phil. Fenn.* 16 (1963) 83–94.

[1965] Semantical analysis of Intuitionistic Logic, I, in *Formal Systems and Recursive Functions*, Crossley et al. eds. North Holland, 1965, pp. 92–130.

Lawson, J.D.

[1979] The duality of continuous posets, *Houston J. Math.* 5 (1979) 357–386.

Linton, F.E.J., and Mikkelsen, C.J.

[1981] Choice as distributivity, manuscript, 1981.

Łoš, J.

[1951] An algebraic proof of completeness for the two-valued propositional calculus, *Coll. Math.* 2 (1951) 236–240.

Łoš, J., and Ryll Nardzewski, C.

[1954] Effectiveness of the representation theory for Boolean algebras, *Fund. Math.* 41 (1954) 49–56.

Markowsky, G.

[1976] Chain-complete posets and directed sets with applications, *Alg. Univ.* 6 (1976) 53–68.

[1981] A motivation and generalization of Scott's notion of a continuous lattice, *Springer Lect. Not. Math.* 871 (1981) 298–307.

Mac Neille, H.M.

[1937] Partially ordered sets, *Trans. Am. Math. Soc.* 42 (1937) 416–460.

McKinsey, J.C.C., and Tarski, A.

[1946] On closed elements in closure algebras, *Ann. Math.* 47 (1946) 122–162.

[1948] Some theorems on the sentential calculi of Lewis and Heyting, *J. Symb. Log.* 13 (1948) 1–15.

Meredith, C.A., and Prior, A.N.

[1963] Notes on the axiomatics of propositional calculus, *Notre Dame J. Form. Log.* 4 (1963)172–187.

Mitschke, G.

[1979] The standardization theorem for the λ-calculus, *Zeit. Math. Log. Grund. Math.* 25 (1979) 29–31.

Monk, J.D., and Bonnet R., eds.

[1989] *Handbook of Boolean algebras*, North Holland, 1989.

Mrowka, S.

[1956] On the ideals extension theorem and its equivalence to the axiom of choice, *Fund. Math.* 43 (1956) 46–49.

Nachbin, L.

[1947] Une propriété caractéristique des algèbres Booléiennes, *Portugal Math.* 6 (1947) 115–118.

[1949] On a characterization of the lattice of all ideals of a Boolean ring, *Fund. Math.* 36 (1949) 137–142.

Nederpelt, R.P.

[1973] *Strong Normalization in a Typed Lambda Calculus with Lambda Structured Types*, Ph.D. Thesis, Eindhoven, 1973.

Nerode, A.

[1957] General topology and partial recursive functionals, *Talks Cornell Summ. Inst. Symb. Log.*, Cornell, 1957, pp. 247–251.

[1990] Some lectures on intuitionistic logic, in *Springer Lect. Not. Math.* 1429 (1990) 12–59.

Nerode, A., and Odifreddi, P.

[1990] *Lambda calculus and constructive logics*, Mathematical Sciences Institute Technical Reports, 55 (1990).

[1994] *Lambda calculus and constructive logics II*, Mathematical Sciences Institute Technical Reports, 44 (1994).

[1997] *Constructive logic from many points of view*, Mathematical Sciences Institute Technical Reports, 4 (1997).

Newman, M.H.A.

[1942] On theories with a combinatorial definition of "equivalence", *Ann. Math.* 43 (1942) 223–243.

Nishimura, I.

[1960] On formulas of one variable in intuitionistic propositional calculus, *J. Symb. Log.* 25 (1960) 327–331.

Odifreddi, P.

[1989] *Classical Recursion Theory*, North Holland, 1989 (Second edition, 1999).

[1989a] *Lecture notes on Lambda Calculus*, Monash University Logic Papers, 65 (1989).

[1990] *Logic and Computer Science* (editor), Academic Press, 1990.

[1997] Short course on Logic, Algebra, and Topology, in *Complexity, Logic, and Recursion Theory*, Sorbi ed., Dekker, 1997, pp. 277–301.

[1999] *Classical Recursion Theory, volume II*, North Holland, 1999.

Ogasawara, T.

[1939] Relation between intuitionistic logic and lattices, *Hirosh. Univ. Sci. Ser.* 9 (1939) 157–164.

Owings, J.C.

[1973] Diagonalization and the recursion theorem, *Notre Dame J. Form. Log.* 14 (1973) 95–99.

Papert, S.

[1959] Which distributive lattices are lattices of closed sets?, *Proc. Camb. Phil. Soc.* 55 (1959) 172–176.

Park, D.M.

[1970] The $\mathcal{Y}$ combinator in Scott's lambda-calculus, *Memo*, University of Warwick, 1970.

Peano, G.

[1884] Addenda to Angelo Genocchi, *Calcolo differenziale e principii di calcolo integrale*, Bocca, 1884.

[1891] Sul concetto di numero, *Rivista di Matematica*, 1 (1891) 87–102 and 256–267.

Peyton Jones, S.L.

[1987] *The Implementation of Functional Programming Languages*, Prentice Hall, 1987.

Plotkin, G.D.

[1972] A set-theoretical definition of application, *Mimeographed Notes*, University of Edinburgh, 1972.

[1980] Lambda definability in the full type hierarchy, in Seldin and Hindley [1980], pp. 363–373.

Pottinger, G.

[1980] A type assignment for the strongly normalizable terms, in *To H.B. Curry: essays on combinatory logics, lambda calculus and formalism*, Seldin et al. eds., Academic Press, 1980, pp. 561–577.

Raney, G.N.

[1952] Completely distributive complete lattices, *Proc. Amer. Math. Soc.* 3 (1952) 677–680.

Rasiowa, H.

[1951] Algebraic treatment of the functional calculi of Heyting and Lewis, *Fund. Math.* 38 (1951) 101–116.

[1974] *An algebraic approach to non-classical logic*, North Holland, 1974.

Rasiowa, H., and Sikorski, R.

[1950] A proof of the completeness theorem of Gödel, *Fund. Math.* 37 (1950) 193–200.

[1963] *The mathematics of metamathematics*, Państwowe Wydawnictwo Naukowe, 1963.

Rieger, L.

[1949] On the lattice theory of Brouwerian propositional logic, *Acta Fac. Rer. Natur. Univ. Carol.* 189 (1949) 1–40.

Rogers, H.

[1959] Computing degrees of unsolvability, *Math. Ann.* 138 (1959) 125–140.

Rosenbloom, P.C.

[1950] *The elements of mathematical logic*, Dover Press, 1950.

Rosser, B.J.

[1935] A mathematical logic without variables, *Ann. Math.* 36 (1935) 127–150.

Sallé, P.

[1978] Une extension de la théorie des types en λ-calcul, *Springer Lect. Not. Comp. Sci.* 62 (1978) 398–410.

Sanchis, L.E.

[1967] Functionals defined by recursion, *Notre Dame J. Form. Log.* 8 (1967) 161–174.

Schönfinkel, M.

[1924] Über die Bausteine der mathematischen Logik, *Math. Ann.* 92 (1924) 305–316, transl. in Van Heijenoort [1967], pp. 357–366.

Schröder, E.

[1891] *Algebra der Logik*, Volume II, 1891.

Schwichtenberg, H.

[1975] Definierbare Funktionen im λ-Kalkül mit Typen, *Arch. Math. Log. Grund. Math.* 17 (1975) 113–114.

Scott, D.S.

[1954] The theorem on maximal ideals in lattices and the axiom of choice, *Bull. Am. Math. Soc.* 60 (1954) 83.

[1954a] Prime ideal theorems for rings, lattices and Boolean algebras, *Bull. Am. Soc.* 60 (1954) 390.

[1963] A system of functional abstraction, *Mimeographed Notes*, 1963.

[1969] Models for the λ-calculus, *Mimeographed Notes*, 1969.

[1970]

[1972] Continuous lattices, *Springer Lect. Not. Math.* 274 (1972) 97–136.

[1975] Lambda Calculus and Recursion Theory, in *Proceedings of the Third Scandinavian Logic Symposium*, Kanger ed., North Holland, 1975, pp. 154–193.

Seldin, J.P., and Hindley, J.R. (editors)

[1980] *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, Academic Press, 1980.

Serre, J.P.

[1955] Faisceaux algébriques cohérents, *Ann. Math.* 61 (1955) 197–278.

Shoenfield, J.R.

[1967] *Mathematical Logic*, Addison Wesley, 1967.

Skolem, T.A.

[1913] *Undersøkelser innenfor logikkens algebra*, Ph.D. Thesis, University of Oslo, 1913.

[1919] Logisch-Kombinatorische Untersuchungen über die Erfüllbarkeit und Beweisbarkeit mathematischen Sätze nebst einem Theoreme über dichte Mengen, *Viden. Krist.* 4 (1919) 1–36.

Statman, R.

[1979] The typed λ-calculus is not elementary recursive, *Theor. Comp. Sci.* 9 (1979) 73–81.

[1980] Completeness, invariance and λ-definability, *J. Symb. Log.* 47 (1980) 17–26.

[1982]

[1985] Logical relations and the typed lambda calculus, *Inf. Contr.* (1985) 85–97.

Stone, M.H.

[1935] Postulates for Boolean algebras and generalized Boolean algebras, *Am. J. Math.* 57 (1935) 703–732.

[1936] The theory of representation for Boolean algebras, *Trans. Am. Math. Soc.* 40 (1936) 37–111.

[1937] Topological representation of distributive lattices and Brouwerian logic, *Čas. Mat. Fys.* 67 (1937) 1–25.

[1937a] Applications of the theory of Boolean rings to general topology, *Trans. Am. Math. Soc.* 41 (1937) 321–364.

[1937b] Algebraic characterization of special Boolean rings, *Fund. Math.* 29 (1937) 223–303.

Tait, W.

[1967] Intensional interpretations of functionals of finite type, I, *J. Symb. Log.* 32 (1967) 198–212.

Takahashi, M.

[1989] Parallel reductions in λ-calculus, *J. Symb. Comp.* 7 (1989) 113–123.

[1995] Parallel reductions in λ-calculus, *Inf. Comp.* 118 (1995) 120–127.

Tarski, A.

[1930] Über einige fundamentale Begriffe der Metamathematik, *Compt. Rend. Soc. Sci. Lett. Varsov.* 23 (1930) 22-29, transl. in [1956], pp. 30–37.

[1935] Zur Grundlegung der Boole'schen Algebra, *Fund. Math.* 24 (1935) 177–198, transl. in [1956], pp. 320–341.

[1935a] Grundzüge der Systemenkalküls, *Fund. Math.* 25 (1935) 503–536, transl. in [1956], pp. 342–383.

[1938] Der Aussasgenkalkül und die Topologie, *Fund. Math.* 31 (1938) 103–134, transl. in [1956], pp. 421–454.

[1954] Prime ideal theorems for Boolean algebras and the axiom of choice, *Bull. Am. Math. Soc.* 60 (1954) 390–391.

[1955] A lattice-theoretical fixed-point theorem and its applications, *Pac. J. Math.* 5 (1955) 285–309.

[1956] *Logic, semantics, metamathematics*, Clarendon Press, 1956.

Turing, A.

[1936] On computable numbers with an application to the Entscheidungsproblem, *Proc. London Math. Soc.* 42 (1936) 230–265, corrections ibidem, 43 (1937) 544-546.

[1937] Computability and λ-definability, *J. Symb. Log.* 2 (1937) 153–163.

[1942] Proof that every typed formula has a normal form, in Seldin and Hindley [1980], pp. 454–455.

Uspenskii, V.A.

[1955] On enumeration operators, *Dokl. Acad. Nauk* 103 (1955) 773–776.

Van Heijenoort, J., ed.

[1967] *From Frege to Gödel*, Harvard University Press, 1967.

Vickers, S.

[1989] *Topology via logic*, Cambridge University Press, 1989.

Wittgenstein, L.

[1921] Logisch-philosophische Abhandlung, *Ann. Naturphil.* 14 (1921) 185–262.

Zermelo, E.

[1908] Untersuchungen über die Grundlagen der Mengenlehre I, *Math. Ann.* 65 (1908) 261–281, transl. in Van Heijenoort [1967], pp. 200–215.

æ